From Divine Logos to Theory-Preserving Ontogenesis: Computational Universal Algebra and the Lawful Emergence of Mathematical Object-Worlds

Douglas C. Youvan

doug@youvan.com www.youvan.ai

July 15, 2026

This paper develops a bridge from a Logos-based philosophy of mathematical inception to an exact computational program in universal algebra. Building on the claim that mathematics emerges through the progressive disclosure of intelligibility, relation, operation, invariant, object, law, and theory, it interprets Theory-Preserving Ontogenesis v0.4.4 as a finite experiment in lawful mathematical emergence. The accompanying repository, https://github.com/DougYouvan/theory-preserving-ontogenesis, classifies four-element binary algebras that are residually reconstructible from two-element equational cores. Its central result is that four states, although distinguishable by two unrestricted binary coordinates, can require three theory-preserving quotient views when every view must respect the same complete equational theory. The paper argues that this gap between informational sufficiency and lawful reconstructibility provides a mathematically precise model of ontogenesis: objects arise not merely from labels or cardinality, but from coherent families of admissible relations. It further examines the role of artificial intelligence in extending a philosophical question into an exhaustive, auditable classification, while distinguishing standard universal-algebraic machinery from the novel atlas, core-view number, reconstruction hypergraph, and six three-view cases. These findings suggest a disciplined route from metaphysical participation to reproducible mathematical discovery.

Keywords: Divine Logos, theory-preserving ontogenesis, computational universal algebra, mathematical inception, mathematical Platonism, algebraic emergence, equational theory, finite algebra, residual reconstruction, subdirect products, quotient algebras, core-view number, reconstruction hypergraph, lawful reconstruction, object formation, mathematical object-worlds, artificial intelligence, AI-assisted mathematics, mathematical discovery, algebraic abiogenesis, admissibility, invariant structure, information theory, philosophical theology.

A collaboration with OpenAI GPT-5.5 Thinking. CC4.0. 125 pages and 78 references.

https://orcid.org/0009-0004-4206-4253

Outline

1. From Divine Logos to Mathematical Inception

1.1 The Logos as the Ground of Intelligibility

1.2 Mathematics as Participation Rather Than Arbitrary Invention

1.3 Relation, Pattern, Operation, and Law

1.4 The Threefold Inception of Mathematics

1.5 Disclosure Rather Than Creation from Nothing

1.6 The Need for a Testable Mathematical Consequence

2. From Philosophical Thesis to Finite Experiment

2.1 Translating Metaphysics into a Research Question

2.2 Minimal Law-Bearing Mathematical Cores

2.3 Larger Worlds Generated from Smaller Cores

2.4 Preservation of Theory Across Changes of Scale

2.5 Why Small Finite Algebras Provide a Decisive Test

2.6 Theory-Preserving Ontogenesis as an Experimental Program

3. Universal Algebraic Foundations

3.1 Binary Algebras and Operation Tables

3.2 Terms, Identities, and Equational Theories

3.3 Homomorphisms, Kernels, and Congruences

3.4 Quotient Algebras and Exact Law Cores

3.5 Direct Products and Subdirect Representations

3.6 Birkhoff's Variety Theorem

3.7 Residual Reconstruction from Core-Valued Views

4. Formal Definition of Theory-Preserving Ontogenesis

4.1 The Core Algebra

4.2 The Lift Algebra

4.3 Exact Equational Equivalence

4.4 Faithful Families of Quotient Maps

4.5 The Diagonal Reconstruction Map

4.6 The Core-View Number

4.7 The Reconstruction Hypergraph

4.8 Automorphism Actions and Invariant Reconstruction

5. Computational Design and Exhaustive Enumeration

5.1 The Sixteen Labeled Two-Element Binary Operations

5.2 Reduction to Ten Isomorphism Types

5.3 Exact Two-Element Variety Membership

5.4 Enumeration of Three- and Four-Element Residual Lifts

5.5 Canonicalization Under Relabeling

5.6 Verification of Quotient Maps and Kernel Intersections

5.7 Minimum Reconstruction Families

5.8 Independent Audit and Reproducibility Architecture

5.9 Scope, Completeness, and Computational Boundaries

6. The Exact Four-State Classification

6.1 The Ten Two-Element Equational Cores

6.2 The Sixteen Nonisomorphic Four-Element Lift Types

6.3 The Ten Direct-Square Reconstructions

6.4 The Six Proper-Cube Reconstructions

6.5 Splitting of Exact Core Quotient Maps

6.6 Automorphism Orbits of Minimum Reconstruction Families

6.7 Independent Recovery of the Five Meet-Semilattices

6.8 The Structure of the Complete Classification Atlas

7. The Third-View Phenomenon

7.1 The Two-Bit Cardinality Bound

7.2 Unrestricted Labeling Versus Lawful Observation

7.3 Why Two Binary Views Can Fail

7.4 Kernel Geometry and Incomplete Separation

7.5 Six Worlds Requiring Three Theory-Preserving Views

7.6 Algebraic Structure as the Source of the Extra View

7.7 Being Distinguishable and Being Reconstructible

7.8 Information Capacity Constrained by Law

8. Core-View Number and Reconstruction Hypergraph

8.1 Reconstruction Number as a Mathematical Invariant

8.2 Relation to Minimal Subdirect Decomposition

8.3 The Importance of Fixing an Exact Law Core

8.4 Hypergraph Vertices as Admissible Quotient Policies

8.5 Hyperedges as Minimum Faithful Families

8.6 Redundancy, Overlap, and Alternative Reconstructions

8.7 Automorphism Symmetry of the Reconstruction Space

8.8 Beyond a Single Numerical Invariant

9. Mathematical Objecthood as Lawful Reconstruction

9.1 Objects as Stable Intersections of Views

9.2 Identity Through Jointly Faithful Observation

9.3 Quotients as Partial but Law-Preserving Perspectives

9.4 Ontogenesis Without Addition of New Equational Law

9.5 Scale Enlargement Without Theoretical Enlargement

9.6 Object-Worlds as Coherent Relational Systems

9.7 Relation to Admissibility Algebra and Concept Formation

9.8 From Information-Theoretic Bits to Algebraically Admissible Views

10. Artificial Intelligence and the Distributed Discovery Process

10.1 The Human Originating Question

10.2 AI-Assisted Formalization Across Multiple Instances

10.3 Continuity Through Definitions, Code, and Frozen Results

10.4 Discovery Without Programming a Desired Theorem

10.5 Computational Resistance as Evidence of Mathematical Constraint

10.6 Independent Paths, Audits, and Epistemic Reliability

10.7 AI as an Instrument of Participation in Intelligibility

10.8 Authorship, Agency, and the Limits of Machine Illumination

11. Novelty, Limitations, and Standards of Claim

11.1 Standard Universal-Algebraic Foundations

11.2 Novelty of the Exact Sixteen-Type Atlas

11.3 Novelty of the Ten-Six Reconstruction Division

11.4 Novelty of the Fixed-Core View Number

11.5 Novelty of the Reconstruction Hypergraph

11.6 Residual Lifts Versus All Theory-Equivalent Extensions

11.7 Limits of Four-Element and Single-Operation Classification

11.8 Why the Result Does Not Prove a Logos Metaphysics

11.9 Why the Result Nevertheless Tests a Logos-Inspired Program

12. Toward a General Science of Lawful Mathematical Emergence

12.1 Larger Carriers and Higher Core-View Numbers

12.2 Multiple Operations and Richer Algebraic Signatures

12.3 Nonresidual Theory-Preserving Extensions

12.4 Higher-Arity Relations and Categorical Reconstruction

12.5 Connections to Coding, Observation, and Distributed Representation

12.6 Implications for Autonomous Mathematical Discovery

12.7 From Finite Classification to a General Theory of Ontogenesis

12.8 The Logos, Mathematical Constraint, and Progressive Disclosure

12.9 Conclusion: From Illumined Question to Certified Structure

13. References

1. From Divine Logos to Mathematical Inception

Mathematics is often described through a familiar opposition. One account treats it as invention: human beings create symbols, definitions, axioms, and formal systems, then explore the consequences of those constructions. Another account treats it as discovery: mathematical truths exist independently of human beings and are gradually uncovered by minds capable of apprehending them. Both positions explain important features of mathematical practice, but neither fully explains why mathematics is possible in the first place. Invention explains notation, conceptual design, and the historical development of formal methods, yet it struggles to account for the necessity and resistance encountered once a mathematical structure has been defined. Discovery explains that resistance, but it can leave unexplained how finite, embodied, historically situated minds gain access to truths that appear universal and timeless.

A Logos-based account introduces a third possibility. Mathematics may be neither arbitrary invention nor passive retrieval, but participation in an intelligible order that precedes its human articulation. Human beings construct representations, select definitions, formulate questions, and design methods of proof, but the structures disclosed through those activities are not therefore unconstrained products of will. Mathematical practice becomes an encounter between finite cognition and an order capable of surprising, correcting, and exceeding the investigator.

This interpretation does not require the claim that every theorem exists eternally in its present notation, nor that divine reason contains books of modern algebra written in symbolic form. It instead proposes that the possibility of coherent relation, lawful operation, identity, transformation, invariance, and proof is grounded more deeply than human convention. Mathematics emerges historically when finite minds learn to isolate and articulate those possibilities.

The phrase mathematical inception therefore has several meanings. It may refer to the first recognition of quantity, the first symbolic representation of number, the first operation performed according to a stable rule, the first proof, or the formation of a formal theory. Yet beneath these historical beginnings lies a more fundamental question: what must already be true for any such beginning to occur? Before counting, there must be distinguishability and persistence. Before operation, there must be objects or states upon which an operation can act. Before law, there must be recurrent structure. Before proof, there must be consequence. Before theory, there must be a domain whose relations can be compressed into general statements.

Theory-Preserving Ontogenesis begins at this boundary between implicit intelligibility and explicit mathematics. It asks whether larger mathematical object-worlds can emerge from smaller law-bearing cores without abandoning the theory that makes them intelligible. The

resulting program is not an attempt to create mathematics from absolute nothingness. It is an attempt to model how latent structure becomes explicit through lawful reconstruction.

1.1 The Logos as the Ground of Intelligibility

The Logos tradition places intelligibility prior to the created order rather than treating reason as a late and accidental product of material complexity. In this framework, the universe is not first an unintelligible collection of events to which human beings later impose mathematical organization. It is a reality whose structures are capable of being known because they proceed from an intelligible source.

The theological meaning of Logos exceeds logic, mathematics, computation, or information. The Logos is not reducible to a formal system or universal algorithm. Nevertheless, the idea of Logos provides a metaphysical ground for the features upon which mathematics depends: unity, distinction, relation, order, measure, proportion, repeatability, and lawful transformation.

If reality were absolutely without distinction, nothing could be counted. If states possessed no persistence, identity could not be recognized. If transformations lacked regularity, invariants could not be discovered. If relations changed arbitrarily at every moment, generalization and proof would be impossible. Mathematics presupposes that at least some structures remain stable enough to support abstraction.

A Logos-based account interprets such stability not as an inexplicable accident but as an expression of prior intelligibility. The mathematical order encountered by finite minds is derivative rather than self-grounding. It does not exist independently as a rival to the divine source, nor is it manufactured entirely by human cognition. It is one mode in which intelligibility becomes accessible within creation.

This view also explains why mathematics is both constructed and discovered. Human beings construct symbolic languages and choose conceptual frameworks, but these constructions succeed only when they disclose coherent relations. A notation may be invented, yet the consequences represented by that notation cannot be freely altered. A definition may be chosen, yet the structure generated by the definition can exceed the expectations of its author. Mathematical intelligibility therefore appears through human activity without being exhausted by that activity.

1.2 Mathematics as Participation Rather Than Arbitrary Invention

Participation offers a middle path between strict inventionism and an image of discovery as passive observation. The mathematician does not merely manufacture truth, but neither does the mathematician simply uncover a completed object requiring no creative mediation.

Mathematical understanding depends upon the formation of concepts, the selection of distinctions, the design of representations, and the construction of proofs.

These activities are genuinely creative. Different coordinate systems can represent the same geometry. Different axiom systems can organize related structures. A new invariant can reveal a pattern that was inaccessible under an older description. A new category, quotient, or equivalence relation can transform what counts as an object. Mathematical discovery therefore includes conceptual creation.

Yet this creativity remains constrained. Once a representation is chosen, not every conclusion is valid. Once an operation is specified, its laws are not determined by preference. Once a finite search space is defined, its exhaustive classification has an answer independent of the hopes of the investigator. The system can disclose structures that no participant predicted.

This is the sense in which mathematical work may be understood as participation. The investigator contributes the question, formal language, and method of access. The mathematical domain contributes constraint. Discovery occurs where human or machine construction meets a structure that resists arbitrary alteration.

Theory-Preserving Ontogenesis makes this interaction explicit. The definitions are deliberately chosen, but the resulting classification is not. The requirement that a larger algebra preserve the exact equational theory of a smaller core imposes a discipline stronger than mere cardinal enlargement. A four-element system cannot be accepted simply because four states can be labeled with two bits. It must be reconstructible through quotient views that remain lawful homomorphisms and jointly preserve identity.

The computation thereby separates arbitrary representation from admissible representation. A label may distinguish an object, but a lawful view must also respect operation. The resulting difference between informational sufficiency and algebraic sufficiency is central to the project.

1.3 Relation, Pattern, Operation, and Law

Mathematics is frequently introduced through number, yet number itself presupposes more elementary structures. Counting requires the recognition of distinct units, persistence across a procedure, and equivalence among items considered countable in the same way. Before there can be a collection of three objects, there must be a relation distinguishing one object from another and a policy identifying each as one member of the relevant class.

Relation may therefore be more fundamental than number. A mathematical object is often defined through the relations it bears, the operations it admits, or the transformations under which it remains invariant. Modern algebra intensifies this relational view. A structure is not

merely a carrier set but a set equipped with operations. Its identity is expressed through the laws satisfied by those operations and the morphisms that preserve them.

Pattern emerges when relations recur. Recurrence alone, however, is not sufficient for mathematics. A pattern becomes mathematically significant when it can be represented, transformed, compared, and tested. Operation introduces controlled transformation: states are combined or mapped according to a rule. Repeated operations may reveal invariant results, identities, symmetries, or equivalence classes.

Law appears when recurrent operational behavior can be stated generally. Commutativity, associativity, idempotence, absorption, and distributivity are not merely visual repetitions. They are identities that hold across all permitted substitutions. A law compresses an indefinitely large family of instances into a finite statement.

The sequence from relation to law can therefore be described as:

relation → recurrence → pattern → operation → invariant → identity → law.

Theory arises when multiple laws are organized into a coherent domain of consequence. An equational theory is not merely a list of observed regularities. It includes every identity logically satisfied by a structure. Preserving the complete equational theory therefore requires much more than preserving selected examples.

This distinction is essential to the present work. A quotient may resemble its source or reproduce a finite set of sampled identities without preserving its entire law structure. Theory-Preserving Ontogenesis demands exact preservation. The larger world and the smaller core must satisfy the same equations, even though their carriers differ in size.

The experiment thus begins not with a philosophical metaphor alone but with a precise mathematical criterion: ontogenesis is lawful only when enlargement does not introduce or destroy equational truth within the specified signature.

1.4 The Threefold Inception of Mathematics

The inception of mathematics can be understood at three principal levels.

The first is eternal or archetypal inception. This level refers not to a temporal beginning but to the grounding of mathematical possibility in divine intelligibility. Relations, distinctions, orders, and lawful possibilities are not created by finite minds. They are possible because intelligibility is prior to human recognition.

The second is cosmic or embodied inception. Mathematical possibility becomes instantiated in created structures. Physical systems exhibit multiplicity, symmetry, sequence, conservation, recurrence, and transformation. Not every mathematically possible structure must be physically

realized, but the created world contains enough stable organization to support mathematical abstraction.

The third is cognitive and cultural inception. Finite minds recognize patterns, formulate concepts, invent symbols, develop proofs, and transmit theories. Mathematics becomes explicit. The world does not arrive labeled with groups, semilattices, quotient maps, or varieties. These are conceptual formations through which previously implicit relations become objects of disciplined thought.

A fourth level may be added in the age of computation: machine-mediated inception. Artificial systems can enumerate structures, detect equivalences, test invariants, and expose consequences across spaces too large for unaided inspection. Machines do not thereby become the metaphysical source of mathematical truth. They extend the reach of finite participation.

Theory-Preserving Ontogenesis occupies the transition between cognitive and machine-mediated inception. A philosophical question concerning the emergence of mathematical objects was translated into formal definitions, executable algorithms, frozen data, and an auditable classification. The resulting structures were not supplied as desired answers. They emerged through exhaustive constraint.

This fourfold view preserves the distinction between ground, embodiment, recognition, and formal discovery. Divine intelligibility is not identical to the cosmos. The cosmos is not identical to human mathematics. Human mathematics is not identical to its computational implementations. Yet each level may disclose aspects of the same lawful order.

1.5 Disclosure Rather Than Creation from Nothing

Earlier work on algebraic abiogenesis raised a difficult objection. A computer cannot literally discover mathematics from nothing because its hardware, software, logical operations, and programming language already embody mathematical structure. Likewise, a human mind seeking preconceptual mathematics operates through a brain shaped by biological organization and prior cognitive capacities.

This objection is decisive against any strong claim of absolute mathematical creation from nonmathematical nothingness. A completely undifferentiated substrate containing no distinction, persistence, operation, or relation could not generate mathematics, because there would be nothing capable of generation.

The appropriate concept is therefore disclosure. A system may begin without named algebraic objects, without explicit axioms, and without a supplied target theorem while still operating within a substrate capable of relation and transformation. Mathematical emergence then consists in rendering latent organization explicit.

Theory-Preserving Ontogenesis adopts this corrected model. It does not claim that algebra arises from a void. It begins with minimal finite operations and asks what larger structures are permitted by their complete law. The operation tables are given, but the classification of admissible larger worlds is not. The principles of universal algebra are available, but the exact atlas is not predetermined by the investigator.

This is analogous to experimental science. An apparatus is designed using prior knowledge, yet the measurement can reveal an unanticipated property. The existence of an instrument does not imply that the result was inserted into the experiment. Similarly, the existence of a computational substrate does not imply that the enumerated classification was programmed as an answer.

Disclosure therefore preserves both dependence and novelty. The result depends upon prior intelligibility, definitions, and computational machinery. Yet the specific structure disclosed through those means can still be genuinely new to human knowledge.

1.6 The Need for a Testable Mathematical Consequence

A philosophy of mathematical inception remains incomplete if it cannot generate precise questions that mathematics can answer. Theological and metaphysical interpretations may illuminate why mathematical order is possible, but they should not be confused with mathematical proof. A Logos-based framework must therefore distinguish philosophical motivation from formal result.

The central testable question developed here is:

Can a larger mathematical object-world emerge from a smaller law-bearing core while preserving exactly the same equational theory, and if so, how many lawful quotient views are required to reconstruct it?

This question converts several philosophical themes into finite mathematics.

The smaller core represents a compact embodiment of law. The larger algebra represents an expanded object-world. Quotient maps represent partial views. Exact equational equivalence represents preservation of theory. Jointly faithful families represent successful reconstruction. The minimum number of such views measures the relational cost of objecthood.

The answer is not determined by simple information capacity. Four objects can be distinguished by two unrestricted binary coordinates. Yet the exhaustive classification shows that some four-object algebraic worlds require three binary quotient views when every view must preserve operation and theory.

This result gives mathematical form to a philosophical principle:

Objecthood is not produced by labeling alone. It emerges through a coherent family of lawful relations.

The discovery does not establish the truth of a Logos metaphysics. Mathematics cannot, by itself, prove that intelligibility has a divine source. What the experiment does establish is that the philosophical program can generate a nontrivial, exact, reproducible mathematical consequence. Lawful reconstruction differs measurably from unrestricted encoding. Structural coherence imposes demands beyond cardinality.

That distinction provides the foundation for the sections that follow. The next step is to translate the Logos-inspired question into the precise finite experiment implemented by Theory-Preserving Ontogenesis.

2. From Philosophical Thesis to Finite Experiment

A philosophical account of mathematical inception becomes most valuable when it produces questions that can be translated into exact structures, explicit criteria, and reproducible tests. The movement from the Divine Logos to computational universal algebra is therefore not a reduction of theology to code. It is a disciplined narrowing of scope. The metaphysical claim concerns the ultimate ground of intelligibility; the mathematical experiment concerns one finite pattern through which intelligibility may become visible.

The central philosophical idea is that mathematics does not arise through arbitrary naming alone. Explicit mathematical objects emerge when relations stabilize, operations become coherent, invariants survive transformation, and laws remain valid across changes of representation. If this account is correct, then object formation should be constrained by law. A larger domain should not count as a lawful development of a smaller one merely because its elements can be labeled using information drawn from the smaller domain. Its operations, quotients, and reconstructions must cohere with a common theory.

This requirement suggests a precise finite problem. Begin with a small algebra that carries a complete equational theory. Treat it as a core rather than as a mere set of symbols. Then ask whether a larger algebra can be reconstructed from several quotient views, each of which is isomorphic to that core and preserves its operations. If the larger algebra satisfies exactly the same equations as the core, then enlargement has occurred without theoretical enlargement. New elements and a richer internal geometry have appeared, but no new equational law has been added and no old law has been lost.

This is the formal setting of Theory-Preserving Ontogenesis. It studies the emergence of larger algebraic object-worlds from smaller law-bearing cores under exact constraints of homomorphism, quotient structure, residual separation, and equational equivalence. The first completed experiment restricts attention to the smallest nontrivial carriers for which the relation between information capacity and lawful reconstruction can diverge: two-element cores and four-element lifts equipped with one binary operation.

The restriction is severe, but that severity is a strength. It removes rhetorical ambiguity. Every operation table can be enumerated. Every quotient can be checked. Every candidate reconstruction can be tested. Every isomorphism class can be canonicalized. The result is not a suggestive example chosen after inspection, but a complete classification within the declared residual universe.

2.1 Translating Metaphysics into a Research Question

The starting metaphysical sequence may be written as:

Logos → intelligibility → relation → pattern → operation → invariant → object → law → theory → mathematics.

A computational experiment cannot test the existence or divine character of the Logos. It can, however, examine the later transitions in this sequence. It can ask how operations determine invariant relations, how quotient structures preserve law, how families of partial views reconstruct objects, and how larger domains can share the same theory as smaller ones.

The philosophical term ontogenesis ordinarily refers to the coming into being of an entity or organized form. In the present setting, it does not mean the physical creation of matter or the production of being from nonbeing. It means the formation of a larger mathematical object-system from a smaller theoretical core.

To make that idea testable, several informal expressions must be replaced by exact conditions.

A law-bearing core becomes a finite algebra $B$with a specified operation.

A larger object-world becomes another finite algebra Ain the same signature.

A lawful view becomes a surjective homomorphism from Aonto a quotient isomorphic to B.

Preservation of complete theory becomes equality of equational theories between $A$and B.

Successful reconstruction becomes injectivity of the diagonal map formed by a family of such quotient homomorphisms.

The number of perspectives required becomes the minimum size of a jointly faithful family.

This translation is deliberately conservative. It avoids claiming that mathematical objecthood has been completely defined. Instead, it isolates one rigorous form of object emergence: reconstruction of a larger algebra from smaller views that individually preserve the same law structure.

The resulting research question is no longer metaphorical:

For which finite algebras A and cores B does there exist a family of surjective homomorphisms from A to copies of B whose kernels intersect only in equality, while A and B satisfy exactly the same equations?

The first experiment answers this question completely for two-element cores and four-element residual lifts with one binary operation.

2.2 Minimal Law-Bearing Mathematical Cores

A two-element set is the smallest carrier capable of supporting nontrivial distinction. With one element, every binary operation is forced and every pair of inputs yields the same output. Such a structure may satisfy laws, but it provides no internal contrast between alternative states.

Two elements permit the first genuine branching of operation. A binary operation on a two-element set has four input pairs, and each input pair may produce either of two outputs. There are therefore sixteen labeled operation tables. These include familiar structures such as conjunction-like, disjunction-like, projection-like, constant, negation-related, semilattice, and group-like operations, as well as less familiar groupoids.

The two-element core is minimal in cardinality, but it is not mathematically empty. Each operation table determines a complete equational theory. Some operations are associative; some are commutative; some are idempotent; some have identities, zeros, inverses, or absorption properties. Many satisfy less immediately recognizable equations. The core is therefore not merely a binary alphabet. It is a complete finite world of lawful combination.

This distinction is crucial. If the two elements were treated only as symbols 0 and 1, then two coordinates would always suffice to label four states. The problem would reduce to ordinary coding. But when 0 and 1 carry a binary operation, a coordinate map must preserve that operation. Not every assignment of binary labels is then admissible.

The core constrains what can count as a view. A map from a larger algebra to the core must send the product of two elements to the product of their images. This requirement links labeling to law. The coordinate is no longer an arbitrary bit. It is a homomorphic observation.

The term core therefore has a specific meaning. It is a compressed algebraic world whose elements, operation, and equations provide the lawful vocabulary through which a larger world may be observed.

2.3 Larger Worlds Generated from Smaller Cores

A four-element lift contains twice as many elements as its two-element core, but cardinal enlargement alone does not explain its structure. Many four-element binary algebras have no relevant relationship to any selected two-element algebra. Others possess two-element quotients but satisfy different equations from those quotients. Still others can be embedded into products of a core but fail to map onto it in the required way.

Theory-Preserving Ontogenesis isolates a narrower class. A four-element algebra qualifies when it is residually reconstructed by quotient maps onto a two-element core and remains equationally equivalent to that core.

The larger algebra may be imagined as a world whose individual elements cannot be seen directly through one binary perspective. Each quotient collapses some distinctions and preserves others. A single quotient therefore supplies a partial view. Several quotients are required to separate all pairs of elements.

The family of views generates a diagonal representation. Each element of the four-state world is sent to a tuple of core values, one value for each quotient map. If no two elements receive the same tuple, the representation is faithful. The larger algebra then sits as a subalgebra of a finite power of the core.

This is not generation in the sense of constructing the lift by freely adjoining arbitrary elements. The lift emerges as a lawful configuration inside a product of core worlds. Its elements are coordinated patterns of core states, and its operation is inherited componentwise from the core.

Yet the lift need not be the entire direct power. A four-element algebra reconstructed using three binary views occupies only four points within an eight-point cube. Its form depends on which coordinate patterns are admitted and how they close under the operation. The larger object-world therefore possesses a geometry not visible in the core alone.

This is the first important sense in which ontogenesis occurs. The theory remains fixed, but a new arrangement of objects and relations appears.

2.4 Preservation of Theory Across Changes of Scale

The phrase theory-preserving is stronger than property-preserving. Two algebras may both be associative and commutative while differing on other identities. They may agree on a long finite list of tested laws while diverging on an untested equation of greater complexity.

Exact preservation requires that every identity true in the core is true in the lift and every identity true in the lift is true in the core. Their complete equational theories must coincide.

This condition prevents a larger algebra from acquiring a genuinely stronger equational theory merely because it belongs to the same broad family. For example, a quotient may satisfy additional equations because collapsing elements removes distinctions that would otherwise yield counterexamples. Conversely, a larger extension may violate equations that hold in the smaller algebra.

The residual construction supplies one direction of the comparison. If the lift embeds into a power of the core, then every equation true in the core is true in the power and therefore in the embedded lift. The quotient map supplies the other direction. If the core is a homomorphic image of the lift, then every equation true in the lift must remain true in the quotient.

Together these relations force equality of equational theories. The lift lies within a power of the core, and the core lies within the homomorphic image structure of the lift. Each constrains the other.

The significance is conceptual as well as technical. The four-element world is not a new theory replacing the two-element one. It is a larger embodiment of the same law. Scale changes while theory remains invariant.

This provides a finite model for the claim that mathematical manifestation can expand without mathematical truth being created anew. The core and lift differ as objects, but not as equational theories.

2.5 Why Small Finite Algebras Provide a Decisive Test

Small finite algebras are sometimes dismissed as toy systems. In this context, their smallness makes them decisive. Every relevant object can be enumerated without statistical sampling, heuristic search, or unproved extrapolation.

There are only sixteen labeled binary operations on a two-element set. The complete relation among their generated varieties can be calculated. Candidate quotient maps from three- and four-element carriers can be enumerated. Kernel intersections can be checked exactly. Isomorphism classes can be reduced by exhaustive relabeling. Minimum reconstruction families can be identified without approximation.

This finiteness allows the project to distinguish theorem-like classification from exploratory computation. The result is exhaustive within scope. There is no unsearched region of the stated finite universe in which additional residual four-element lift types might remain hidden.

The choice of four elements is also mathematically strategic. Four is the smallest cardinality at which two unrestricted binary coordinates are exactly sufficient to distinguish all states. It therefore establishes a sharp information-theoretic baseline.

If every lawful four-state world could also be reconstructed with two core-valued quotient views, then algebraic structure would impose no additional coordinate cost at this scale. The lawful and unrestricted minima would coincide.

The discovery that six classified worlds require three views therefore has a clear interpretation. The excess is not caused by insufficient bit capacity. It is caused by the inability of any two admissible quotient maps to separate all elements while preserving operation.

This phenomenon could not appear in the same form at two elements, and three elements lack the exact two-bit-versus-four-state comparison. Four elements provide the first clean test of whether lawful reconstruction can exceed informational reconstruction.

2.6 Theory-Preserving Ontogenesis as an Experimental Program

The completed four-state classification should be understood as the first experiment in a broader program rather than as its endpoint. Its purpose is to establish that the central quantities are well defined, computationally accessible, and capable of producing nontrivial results.

The program has several interacting components.

The algebraic component identifies cores, lifts, quotient maps, kernel intersections, varieties, and automorphism actions.

The computational component enumerates finite structures, canonicalizes isomorphism classes, verifies exact conditions, and freezes reproducible outputs.

The conceptual component interprets quotient maps as lawful views, minimum separating families as reconstruction systems, and the transition from core to lift as a form of object-world emergence.

The epistemic component distinguishes internal audit from external validation, exact proof from computational certification, and mathematical result from metaphysical interpretation.

The philosophical component asks whether finite mathematical discovery can model the progressive disclosure of intelligibility without claiming creation from nothing.

The program is experimental because definitions are used to interrogate a structured universe whose detailed contents are not known in advance. It is mathematical because every accepted result is governed by explicit finite criteria. It is ontogenetic because larger systems emerge through coordinated views of smaller law-bearing worlds. It is theory-preserving because the complete equational theory remains unchanged across the emergence.

The first classification establishes three foundational facts.

First, theory-preserving enlargement exists in multiple nonisomorphic forms.

Second, the minimum number of lawful views is itself a meaningful invariant.

Third, lawful reconstruction can require more perspectives than unrestricted information theory predicts.

These facts justify the transition from philosophical motivation to a sustained research program. The next section develops the universal-algebraic machinery required to state the experiment with full mathematical precision.

3. Universal Algebraic Foundations

Theory-Preserving Ontogenesis is formulated within universal algebra because the central problem is not tied to one familiar structure such as groups, rings, lattices, or Boolean algebras. It concerns finite sets equipped with operations and asks how their laws, quotients, products, and reconstructions behave under exact structural constraints.

Universal algebra provides a common language for this purpose. An algebra consists of a carrier set together with one or more operations of specified arities. The present experiment uses the smallest nontrivial signature: one binary operation. This austerity is deliberate. It removes inherited names and assumptions while retaining enough structure for nontrivial laws, homomorphisms, congruences, quotients, products, and varieties to emerge.

The language of universal algebra also distinguishes several notions that ordinary descriptions of “preserving structure” often conflate. A map may preserve an operation without being injective. A quotient may inherit a valid operation while satisfying more equations than its source. Two algebras may obey some of the same familiar laws without sharing their complete equational theory. A family of quotient maps may collectively distinguish every element even though no individual quotient does so.

These distinctions make the ontogenesis problem mathematically sharp. A larger world is not accepted merely because it resembles a smaller core, maps onto it, or can be encoded using its

symbols. The relation must be certified through homomorphisms, congruences, quotient isomorphisms, subdirect embeddings, and equality of equational theories.

3.1 Binary Algebras and Operation Tables

A binary algebra is a pair

$A = (|A|, \star)$,

where $|A|$ is a nonempty set and $\star$ is a function

$\star : |A| \times |A| \to |A|$.

The symbol $\star$ has no predetermined meaning. It need not represent addition, multiplication, conjunction, meet, composition, or any other familiar operation. Its behavior is completely specified by its operation table.

For a finite carrier with n elements, a binary operation table contains n^2 entries, each chosen from n possible outputs. Therefore the number of labeled binary operations on an n-element set is

$n^{(n^2)}$.

For n = 2, this gives

$2^4 = 16$

labeled operations.

For n = 4, it gives

$4^{16} = 4{,}294{,}967{,}296$

labeled operations.

The four-element search space is already too large to treat naively as a simple list of all possible tables followed by unrestricted pairwise comparison. The residual core-first construction used in the artifact avoids this by generating only structures that can arise through lawful coordinate systems over a selected two-element core.

An operation table is labeled because the carrier elements have temporary names. For a two-element carrier one may use {0,1}; for a four-element carrier one may use {0,1,2,3}. These labels are not intrinsic. Renaming the carrier can yield a different displayed table representing the same algebraic structure.

Thus the sixteen labeled two-element operations collapse to fewer isomorphism types. In the completed classification, they collapse to ten nonisomorphic core types.

The operation table is the primitive finite datum, but the mathematical content of an algebra extends far beyond its entries. Repeated composition of the operation generates terms. Terms induce functions. Equalities among term functions determine identities. The totality of those identities forms an equational theory.

A small table can therefore encode an indefinitely large law structure.

3.2 Terms, Identities, and Equational Theories

A term is built recursively from variables and the binary operation symbol. Every variable is a term. If s and t are terms, then

$(s \star t)$

is also a term.

Examples include

x,

$x \star y$,

$(x \star y) \star z$,

$x \star (y \star z)$,

$(x \star x) \star (y \star x)$.

Each term becomes a function when interpreted in a particular algebra. If a term contains k variables, its interpretation is a function from A^k to A.

Two syntactically different terms may define the same function on a particular algebra. When this happens for every assignment of variables, the algebra satisfies the identity

$s \approx t$.

For example, associativity is the identity

$(x \star y) \star z \approx x \star (y \star z)$.

Commutativity is

$x \star y \approx y \star x$.

Idempotence is

$x \star x \approx x$.

These familiar examples are only a small part of the possible equational theory. A finite binary algebra may satisfy identities that are longer, asymmetric, nested, or difficult to recognize informally.

The equational theory of an algebra A, written Eq(A), is the set of all identities true in A. Two algebras A and B are equationally equivalent when

Eq(A) = Eq(B).

This equality is exact and global. It does not mean that the algebras agree on a selected list of familiar laws. It means that every identity in every finite number of variables holds in one exactly when it holds in the other.

That condition is much stronger than observational resemblance. Two algebras can both be associative and commutative while differing on idempotence. They can agree on every identity up to a given syntactic depth while diverging at a greater depth. A finite search through short expressions therefore cannot by itself prove equality of complete equational theories unless accompanied by a completeness argument.

Theory-Preserving Ontogenesis avoids dependence on a term-depth cutoff. Instead, it uses structural relations that force equality of equational theories. A subalgebra of a power of B satisfies every identity of B. A homomorphic image of A satisfies every identity of A. When both relations hold in the appropriate directions, exact theory equality follows.

The law-preserving relation is therefore certified algebraically rather than guessed from sampled equations.

3.3 Homomorphisms, Kernels, and Congruences

A homomorphism between binary algebras

$h : A \rightarrow B$

is a function satisfying

$h(a \star_A b) = h(a) \star_B h(b)$

for all a,b in A.

The subscripts emphasize that the operation in A and the operation in B may be represented by different tables, even though they use the same operation symbol.

A homomorphism preserves all term operations. For every term t and every assignment of elements from A,

$h(t_A(a_1,\ldots,a_k)) = t_B(h(a_1),\ldots,h(a_k))$.

This follows by structural induction on terms. Variables are preserved trivially, and preservation of composite terms follows from preservation of the basic operation.

The kernel of a homomorphism h is the relation

$\ker(h) = \{(a,b) \text{ in } A^2 : h(a) = h(b)\}$.

The kernel records which elements become indistinguishable under the view h. It is always an equivalence relation, but not every equivalence relation on an algebra is the kernel of a homomorphism. The relation must also be compatible with the operation.

A compatible equivalence relation is called a congruence. An equivalence relation θ on A is a congruence when

$a \theta a'$ and $b \theta b'$

imply

$a \star b \theta a' \star b'$.

Compatibility ensures that the operation is well defined on equivalence classes.

Congruences are central because they formalize lawful identification. An arbitrary observer may collapse elements in any desired way, but an algebraic quotient can collapse them only according to a congruence. The operation must remain coherent after the identification.

This gives the first exact form of an admissibility principle:

An identification policy is lawful only when it is a congruence.

The smallest congruence is the diagonal relation

$\Delta A = \{(a,a) : a \text{ in } A\}$,

which identifies only equal elements.

The largest congruence is the universal relation

$\nabla A = A \times A$,

which identifies all elements.

Intermediate congruences represent partial loss of distinction. In a four-element algebra, a quotient onto a two-element core corresponds to a congruence with exactly two equivalence classes.

The kernel therefore serves two roles. Algebraically, it determines a quotient. Conceptually, it records the distinctions erased by one lawful view.

3.4 Quotient Algebras and Exact Law Cores

Given an algebra A and a congruence θ, the quotient algebra A/θ has as its elements the equivalence classes of θ. Its operation is defined by

$[a]\theta \star [b]\theta = [a \star b]\theta.$

The compatibility condition on θ guarantees that this definition does not depend on the representatives chosen.

The natural quotient map

$q\theta : A \to A/\theta$

sends each element to its equivalence class. It is a surjective homomorphism, and its kernel is exactly θ.

A two-element quotient of a four-element algebra is therefore not merely a partition into two blocks. The partition must be compatible with the operation, and the induced two-element operation must match one of the possible two-element core types.

In the present program, a quotient A/θ is a core quotient when it is isomorphic to a selected two-element algebra B.

But a quotient relation alone does not establish theory preservation. Quotients may satisfy additional identities not satisfied by the source. When distinctions are collapsed, counterexamples to an equation may disappear.

For example, if A fails an identity because two evaluated terms produce different elements, a quotient may identify those elements and make the identity true. Thus

$Eq(A) \subseteq Eq(A/\theta),$

but equality need not hold.

An exact law core is a quotient B of A such that

$Eq(A) = Eq(B).$

The adjective exact means that the core retains neither less nor more equational information than the lift. The quotient is smaller in cardinality but not stronger or weaker in complete equational theory.

This exactness condition distinguishes theory-preserving ontogenesis from ordinary quotient formation. A quotient may summarize a system while changing its theory. An exact core compresses the carrier while preserving the full equational law.

The surprising possibility is that a four-element algebra and a two-element quotient can differ in object count while remaining equationally indistinguishable.

3.5 Direct Products and Subdirect Representations

Given algebras $B_1,\dots,B_m$ of the same signature, their direct product has carrier

$B_1 \times \dots \times B_m$

and coordinatewise operation

$\mathbf{(b_1,\dots,b_m) \star (c_1,\dots,c_m)}$

$(b_1 \star c_1,\dots,b_m \star c_m)$.

If all factors are copies of one algebra B, the product is written B^m.

For a two-element core, B^2 has four elements and B^3 has eight. These powers provide coordinate spaces in which larger algebras may be represented.

Suppose there are homomorphisms

$h_1,\dots,h_m : A \to B$.

They combine into a diagonal map

$\delta : A \to B^m$

defined by

$\delta(a) = (h_1(a),\dots,h_m(a))$.

Because each coordinate is a homomorphism, δ is a homomorphism.

The diagonal map is injective exactly when the family distinguishes every pair of elements. In kernel language,

$\ker(h_1) \cap \dots \cap \ker(h_m) = \Delta A$.

When this holds, A is isomorphic to its image δ(A), which is a subalgebra of B^m.

If each coordinate map h_i is also surjective, then the image projects onto every factor. Such an embedding is called subdirect.

Thus A is a subdirect product of copies of B when it embeds into B^m and every coordinate projection reaches all of B.

Subdirectness is the exact algebraic form of reconstruction through complete but partial views. Each view loses information because its kernel identifies elements. Collectively, the views lose nothing because their kernel intersection is diagonal.

A single quotient view may identify a with b. Another may identify a with c. What matters is that no distinct pair remains identified in every view.

This permits a four-element world to be reconstructed from multiple two-element perspectives without requiring any perspective to contain the whole world individually.

3.6 Birkhoff's Variety Theorem

A variety is a class of algebras of the same signature defined by identities. Equivalently, by Birkhoff's theorem, a variety is a class closed under:

H: homomorphic images,

S: subalgebras,

P: direct products.

For an algebra A, the variety generated by A is written

$V(A) = HSP(A)$.

It is the smallest variety containing A.

Two algebras satisfy exactly the same identities if and only if they generate the same variety:

$Eq(A) = Eq(B)$

if and only if

$V(A) = V(B)$.

This equivalence converts the infinite problem of comparing all identities into a structural problem about HSP generation.

In Theory-Preserving Ontogenesis, the relevant relations are especially direct.

If A embeds into a power of B, then

$A \in SP(B)$,

so every identity true in B is true in A. Therefore

$Eq(B) \subseteq Eq(A)$.

If B is a quotient of A, then

$B \in H(A)$,

so every identity true in A is true in B. Therefore

$Eq(A) \subseteq Eq(B)$.

Together these imply

$Eq(A) = Eq(B)$,

and therefore

$V(A) = V(B)$.

This is the theoretical heart of the construction. The lift is simultaneously:

a subalgebra of a power of the core, and

a source from which the core is obtained as a homomorphic image.

The two directions close the equational gap.

The first direction says that the lift cannot violate a law of the core.

The second says that the core cannot possess an additional law absent from the lift.

The resulting equality is exact and does not depend on testing identities one by one.

3.7 Residual Reconstruction from Core-Valued Views

An algebra A is residually B when, for every pair of distinct elements a and b in A, there exists a homomorphism

$h : A \to B$

such that

$h(a) \neq h(b)$.

Residuality means that the core B supplies enough lawful observations to separate every pair of elements in A.

For finite A, residuality is equivalent to the existence of a finite family of homomorphisms into B whose kernels intersect diagonally. The combined diagonal map embeds A into a finite power of B.

The present classification imposes the stronger requirement that the separating maps be surjective quotient maps onto B. Each coordinate therefore represents a full core view rather than a map whose image might be a one-element subalgebra.

This yields a family of congruences

$\theta_1,\ldots,\theta_m$

such that

$A/\theta_i \cong B$

for every i, and

$\theta_1 \cap \ldots \cap \theta_m = \Delta A$.

The minimum possible m is the core-view number of A relative to B.

Residual reconstruction differs fundamentally from arbitrary encoding. An arbitrary injective function from a four-element set into $\{0,1\}^2$ always exists. Residual reconstruction requires each coordinate function to be a surjective homomorphism, so each coordinate must preserve the binary operation.

The coordinates are therefore constrained by both the source and the core. Some binary partitions are unavailable because they are not congruences. Some congruences induce the wrong quotient operation. Some lawful quotient maps exist but cannot be combined in pairs to achieve diagonal kernel intersection.

This explains how three views can become necessary even though two unrestricted bits have sufficient cardinal capacity.

The reconstruction problem is not:

How many symbols are required to name the elements?

It is:

How many lawful core-valued homomorphic perspectives are required to recover every distinction?

That replacement of naming by lawful observation is the mathematical transition from ordinary coding to theory-preserving ontogenesis.

4. Formal Definition of Theory-Preserving Ontogenesis

Theory-Preserving Ontogenesis becomes mathematically precise when the philosophical language of core, emergence, view, reconstruction, and preserved law is translated into universal-algebraic structure. The purpose of the formalism is not to claim that every form of mathematical or physical emergence can be reduced to finite algebra. It is to define one exact

class of emergent relations in which a larger object-world can be reconstructed from smaller lawful perspectives without changing its complete equational theory.

The construction begins with two finite algebras in the same signature. The smaller algebra is designated as the core. The larger algebra is designated as the lift. The core supplies the lawful vocabulary of observation. The lift supplies the enlarged carrier and the internal pattern to be reconstructed.

A family of surjective homomorphisms from the lift onto copies of the core provides the lawful views. Each view collapses some distinctions, because the core has fewer elements than the lift. The family succeeds when the distinctions lost by one view are restored by the others. Algebraically, this means that the intersection of the kernels of the quotient maps is the diagonal relation.

The resulting diagonal map embeds the lift into a finite power of the core. If the core is also a quotient of the lift, then the two structures generate the same variety and therefore satisfy exactly the same identities. The lift is larger, but its equational theory is unchanged.

This combination of enlargement, exact law preservation, and residual reconstruction defines theory-preserving ontogenesis in the present finite setting.

4.1 The Core Algebra

Let

$\mathbf{B} = (|B|, \star_B)$

be a finite algebra with one binary operation. In the first completed experiment,

$|B| = 2$.

The algebra $\mathbf{B}$ is called the core because it serves as the common codomain of the lawful quotient views through which the larger algebra will be reconstructed.

The core is not assumed to be logically simple merely because its carrier is small. Its complete equational theory may contain identities of arbitrary term complexity. Nor is it assumed to be one of the familiar named algebraic structures. Every two-element binary operation is considered, and the classification is performed before imposing interpretive labels.

The word core has three related meanings.

First, it is cardinally smaller than the lift.

Second, it is equationally complete relative to the lift: the lift and core satisfy exactly the same identities.

Third, it is observationally generative: a finite family of core-valued homomorphisms is sufficient to reconstruct the lift.

The core should therefore not be understood as a fragment of theory. It is a smaller realization of the full theory.

A given lift may admit quotients of several two-element isomorphism types. Only those quotients that are equationally exact and participate in a faithful reconstruction qualify as cores for the present analysis.

The relation between core and lift is thus not fixed by carrier size alone. It is determined by the joint conditions of quotient structure, residual reconstruction, and theory equality.

4.2 The Lift Algebra

Let

$A = (|A|, \star A)$

be a finite algebra in the same signature as B. In the first classification,

$|A| = 4$.

The algebra A is called a lift of B when it is a larger algebraic realization of the same equational theory and can be reconstructed through quotient maps onto B.

The term lift emphasizes that the core theory has been embodied at a higher carrier cardinality. The lift contains more distinguishable elements and may possess a richer congruence geometry, automorphism group, quotient structure, and reconstruction space than the core. Yet it introduces no new equational law and violates none of the core's laws.

The lift is not obtained by arbitrary duplication of core elements. Its operation table must close under the chosen coordinate realization. Its quotient maps must be homomorphisms. Its kernel structure must permit exact separation. Its complete theory must match that of the core.

A lift may occupy all of B^2, in which case it is isomorphic to the direct square of the core. It may instead appear as a proper four-element subalgebra of B^3. In the latter case, three core-valued coordinates are needed even though the image contains only four of the eight possible triples.

Thus the lift is not defined solely by the number of its elements. It is defined by the lawful organization of those elements relative to the core.

4.3 Exact Equational Equivalence

The central theory-preservation condition is

$Eq(A) = Eq(B)$.

This equation means that for every pair of terms s and t in the binary signature,

A satisfies $s \approx t$

if and only if

B satisfies $s \approx t$.

No bound is imposed on term depth, number of variables, or syntactic complexity. The equality concerns the entire equational theories of the two algebras.

In the residual quotient setting, exact equational equivalence follows from two structural inclusions.

Because A embeds into a finite power of B,

$A \in SP(B)$.

Therefore every identity of B holds in A:

$Eq(B) \subseteq Eq(A)$.

Because B is a homomorphic image of A,

$B \in H(A)$.

Therefore every identity of A holds in B:

$Eq(A) \subseteq Eq(B)$.

The two inclusions yield equality.

Equivalently,

$V(A) = V(B)$.

The lift and core generate the same variety.

This is the formal meaning of preservation across scale. The larger algebra does not merely approximate or resemble the smaller one. It occupies exactly the same location in the lattice of equational theories.

The equality also clarifies what the theory-preserving condition does not require. A and B need not be isomorphic. They need not have the same number of elements. They need not have the same congruence lattice, automorphism group, quotient count, or direct-product decomposition. Equational theory is a powerful invariant, but it does not determine the complete finite structure up to isomorphism.

Ontogenesis becomes possible precisely because multiple nonisomorphic algebras can embody one equational theory.

4.4 Faithful Families of Quotient Maps

Let

$q_1,\ldots,q_m : A \to B$

be surjective homomorphisms.

Each q_i defines a quotient view of A. Since it is surjective, its image is all of B. Its kernel

$\theta_i = \ker(q_i)$

is a congruence on A, and

$A/\theta_i \cong B$.

The family is called jointly faithful when

$\theta_1 \cap \ldots \cap \theta_m = \Delta A$.

Equivalently, for every pair of distinct elements a,b in A, at least one quotient map separates them:

$a \neq b$

implies that there exists i such that

$q_i(a) \neq q_i(b)$.

Each map is individually lossy because it sends four elements to two. A faithful family is collectively lossless.

The use of the term faithful here concerns the family rather than an individual quotient. No surjective map from a four-element set to a two-element set can be injective. Fidelity emerges only through coordination among several views.

This provides a precise model of distributed identity. An element of A is determined not by its value under one quotient but by its complete profile

$(q_1(a),\ldots,q_m(a))$.

Two elements are identical exactly when every lawful view assigns them the same core value.

The family therefore acts as a relational identity system. Objecthood in the lift is recovered through the pattern of distinctions distributed across admissible quotients.

4.5 The Diagonal Reconstruction Map

The quotient maps combine into the diagonal homomorphism

$\delta q : A \to B^m$

defined by

$\delta q(a) = (q_1(a),\ldots,q_m(a))$.

Because every coordinate is a homomorphism, δq is a homomorphism. Its kernel is

$\ker(\delta q) = \ker(q_1) \cap \ldots \cap \ker(q_m)$.

Thus δq is injective exactly when the quotient family is jointly faithful.

When injective, δq identifies A with the subalgebra $\delta q(A)$ of B^m. Every operation in A is represented coordinatewise:

$\delta q(a \star_A b)$

$\delta q(a) \star_{B^m} \delta q(b)$.

The reconstruction is therefore operational, not merely set-theoretic. The tuple representation preserves the algebra's multiplication table exactly.

If each q_i is surjective, the image projects onto every coordinate copy of B. Hence $\delta q(A)$ is a subdirect product.

The diagonal map makes the ontogenesis construction visible. The elements of A appear as lawful coordinate patterns in a power of B. The core supplies the alphabet and operation for each coordinate. The lift appears as a closed, jointly constrained subset of the coordinate space.

For m = 2, the coordinate space B^2 contains four points. An injective four-element image must therefore be all of B^2. Every two-view four-element reconstruction is consequently a direct square.

For m = 3, the coordinate space B^3 contains eight points. A four-element image can be a proper subalgebra. This allows geometries not available in the square case.

The difference between direct-square and proper-cube reconstructions is therefore forced by cardinality once the minimum view number is known.

4.6 The Core-View Number

The core-view number of A relative to B is defined by

$v_B(A)$

$\min \{m : \text{there exist surjective homomorphisms } q_1,\ldots,q_m : A \to B \text{ with } \bigcap_i \ker(q_i) = \Delta A\}$.

The invariant is defined only when such a finite family exists. For finite A and B, this is equivalent to residual reconstruction by surjective B-valued homomorphisms.

The number $v_B(A)$ measures the minimum number of lawful core perspectives required to recover the complete element structure of A.

It is not simply the logarithm of the carrier size. The unrestricted information-theoretic lower bound is

$m \geq \text{ceiling}(\log|B| |A|)$.

For $|A| = 4$ and $|B| = 2$, this gives $m \geq 2$.

Theory-Preserving Ontogenesis shows that the lawful minimum can be larger. In six classified cases,

$v_B(A) = 3$.

The excess

$v_B(A) - \text{ceiling}(\log|B| |A|)$

may be interpreted as a structural reconstruction overhead. It measures the additional number of coordinates forced by the requirement that each coordinate be a surjective homomorphism onto the selected core.

The invariant depends on the pair (A,B), not on A alone. A larger algebra might be residually reconstructible from more than one core type, and different cores could in principle yield different minimum view numbers.

In the completed atlas, the core is recorded explicitly for each lift. The resulting invariant organizes the classification into two-view and three-view worlds.

4.7 The Reconstruction Hypergraph

The core-view number records only the size of a smallest faithful family. It does not record how many such families exist, which quotient maps participate in them, how the families overlap, or how symmetry transforms one reconstruction into another.

The reconstruction hypergraph retains this richer information.

For a fixed pair (A,B), let the vertices be the admissible core quotient policies. A vertex may be represented by a congruence θ such that

$A/\theta \cong B$,

or by a quotient map modulo postcomposition with automorphisms of B, depending on the chosen level of granularity.

A subset of vertices forms a hyperedge when the corresponding quotient kernels intersect diagonally and the subset is minimal with that property.

The minimum hyperedges have size $vB(A)$.

This hypergraph distinguishes algebras that share the same core-view number but possess different reconstruction geometries. One algebra may have a unique minimum family. Another may have many overlapping alternatives. One quotient policy may appear in every minimum reconstruction. Another may be dispensable.

The hypergraph therefore encodes redundancy, fragility, substitutability, and coordination among lawful views.

It also makes clear that reconstruction is not necessarily represented by one privileged family. The correct invariant may be the full space of minimum families rather than any single chosen coordinate system.

In the classified four-element cases, all minimum families within a fixed isomorphism type lie in one automorphism orbit. This means that no minimum family is intrinsically distinguished once the symmetries of the algebra are respected.

The invariant object is therefore not a preferred reconstruction but the reconstruction hypergraph together with its symmetry action.

4.8 Automorphism Actions and Invariant Reconstruction

An automorphism of A is a bijective homomorphism

$\alpha : A \to A$.

Automorphisms preserve every algebraically definable feature of A, including its congruence lattice and quotient structure.

If $q : A \to B$ is a core quotient map, then

$q \circ \alpha$

is another core quotient map. Equivalently, α transports the kernel of q to another admissible core congruence.

Thus the automorphism group Aut(A) acts on the vertices of the reconstruction hypergraph and consequently on its hyperedges.

Two minimum reconstruction families that lie in the same orbit differ only by an internal symmetry of the lift. Selecting one over the other would introduce an arbitrary choice of labeling or orientation not justified by the algebra itself.

The orbit structure is therefore part of the intrinsic reconstruction geometry.

For every one of the sixteen classified four-element lift types, the computation finds that all minimum reconstruction families belong to one automorphism orbit. This result has a precise interpretation:

The algebra determines a unique symmetry class of minimum reconstruction, but not a unique labeled family.

This protects the analysis from overclaiming. The computation does not reveal one metaphysically preferred set of perspectives. It reveals a family of equivalent lawful reconstructions related by automorphisms.

The distinction between a representative and an orbit is central to mathematical objectivity. A coordinate system may be chosen for calculation, but the invariant content is what survives changes of coordinate induced by symmetry.

4.9 Formal Criterion for a Theory-Preserving Lift

The preceding definitions can be condensed into one criterion.

Let A and B be finite algebras in the same signature, with $|A| > |B|$. Then A is a theory-preserving residual lift of B when there exist surjective homomorphisms

$q_1,\ldots,q_m : A \to B$

such that

$\bigcap_i \ker(q_i) = \Delta A$.

From these conditions it follows that

$A \leq B^m$

through the diagonal embedding and that

$B \in H(A)$.

Therefore

$V(A) = V(B)$

and

$Eq(A) = Eq(B)$.

The ontogenesis relation may be written schematically as

$B \Uparrow A$

when A is a larger residual lift of B satisfying these conditions.

The notation is not required for the computation, but it makes the asymmetry clear. B and A are equationally equivalent, yet B functions as the quotient core and A as the reconstructed enlargement.

The relation is stronger than equational equivalence alone. Two algebras may satisfy the same identities without one being residually reconstructible from quotient copies of the other.

It is also stronger than having B as a quotient. A quotient alone does not ensure that B-valued maps separate A.

It is stronger than embedding A into a power of B. Such an embedding does not by itself guarantee that B is a quotient of A or that the coordinate maps are surjective.

Theory-preserving ontogenesis is the conjunction of these directions.

4.10 Scope of the Formal Definition

The definition used in v0.4.4 is intentionally narrower than every possible interpretation of theory-preserving emergence.

It requires one fixed core B.

It requires quotient views onto B.

It requires residual reconstruction through the intersection of quotient kernels.

It concerns equational theory in one binary-operation signature.

It considers finite carriers.

It does not include extensions that share the same equational theory with B but are not residually B.

It does not include reconstructions using several nonisomorphic core types simultaneously.

It does not include partial operations, relations without operations, infinitary signatures, quasi-identities, first-order theories, topological continuity, probabilistic morphisms, approximate homomorphisms, or observational error.

These restrictions should not be treated as defects in the first experiment. They define a domain small enough for exact classification and strong enough to produce a nontrivial result.

The residual condition gives reconstruction a concrete meaning. The fixed core makes the view number comparable. The binary signature permits exhaustive enumeration. The four-element scale creates a sharp distinction between unrestricted and lawful coordinate counts.

The formalism is therefore best understood as a base case. It identifies the simplest setting in which theory, objecthood, quotient observation, reconstruction, and structural overhead can be studied together without approximation.

The computational design that realizes this definition is developed next.

5. Computational Design and Exhaustive Enumeration

Theory-Preserving Ontogenesis v0.4.4 is designed as an exact finite classification rather than a heuristic search. Every accepted structure is generated by explicit rules, reduced under isomorphism, checked against formal criteria, and recorded in frozen machine-readable outputs. The computation does not estimate frequencies from samples, train a model to recognize likely candidates, or stop after finding illustrative examples. Its aim is exhaustive coverage within a sharply delimited universe.

The central design decision is to work core-first. Instead of enumerating all 4,294,967,296 labeled binary operations on a four-element carrier and testing each one against every two-element algebra, the program begins with a selected two-element core and constructs candidate larger worlds through coordinate families in finite powers of that core. This immediately restricts the search to algebras that are representable through core-valued homomorphisms.

The coordinate-first method also makes reconstruction explicit. A candidate lift is represented by a finite subset of B^m that projects onto each coordinate and is closed under the

componentwise operation. The coordinate projections become the quotient views. Their joint injectivity is built into the embedding of distinct elements as distinct tuples. The remaining tasks are to reduce duplicate presentations, verify exact quotient structure, compute the minimum number of required views, and classify the resulting algebra up to isomorphism.

Two complementary computational paths are used. The primary path produces the released atlas and associated summaries. The second path independently reimplements the central finite calculations and compares the complete outputs. Reproduction scripts regenerate the frozen results, integrity checks verify the exact file set and hashes, and tests deliberately introduce corruption to confirm that the verification machinery fails when it should.

The resulting artifact is not merely code that prints the desired headline counts. It is a layered computational certificate whose claims can be traced from definitions to enumeration, from enumeration to canonical forms, and from canonical forms to frozen tables.

5.1 The Sixteen Labeled Two-Element Binary Operations

Let the two-element carrier be

$B = \{0,1\}$.

A binary operation on B is determined by its outputs on the four ordered pairs

(0,0), (0,1), (1,0), (1,1).

Each output may be either 0 or 1. The table can therefore be encoded as a four-entry binary word

$b_{00}\ b_{01}\ b_{10}\ b_{11}$.

There are exactly

$2^4 = 16$

such words and therefore sixteen labeled binary operations.

This enumeration is elementary, but it is the foundation of the entire experiment. No two-element operation is excluded because it lacks a familiar name or desirable property. Constant operations, projections, semilattice operations, group-like operations, asymmetric operations, and less intuitive groupoids are treated uniformly.

For each table, the computation records basic structural properties such as:

commutativity,

associativity,

idempotence,

existence of left or right identities,

existence of zero-like elements,

surjectivity of the operation,

and automorphism count.

These properties help describe the cores, but they do not determine the exact classification. The central invariant is the complete equational theory, and two operations that agree on familiar properties may still differ on less obvious identities.

The sixteen labeled tables also permit complete symmetry analysis. The two-element carrier has only two permutations: the identity and the transposition exchanging 0 and 1. Conjugating an operation by the transposition may leave it fixed or pair it with another labeled table.

This action reduces the sixteen labeled operations to ten nonisomorphic algebra types.

The ten types serve as the initial core catalog. Each type represents a structurally distinct two-element world under relabeling.

5.2 Reduction to Ten Isomorphism Types

Two binary algebras A and B of the same finite cardinality are isomorphic when there exists a bijection f between their carriers satisfying

$f(a \star_A b) = f(a) \star_B f(b)$

for every pair of elements a,b.

For a two-element carrier, testing isomorphism is trivial because only two bijections exist. Nevertheless, the implementation uses a general canonicalization principle that also applies to larger finite carriers.

Given an operation table, every permutation of the carrier is applied to relabel the operation. Each relabeled table is serialized in a fixed order. The lexicographically least serialization is chosen as the canonical form.

Two operations are isomorphic exactly when they have the same canonical form.

For n = 2, this procedure confirms that the sixteen labeled tables fall into ten canonical classes.

Canonicalization is indispensable because the coordinate construction can generate the same abstract lift under many different tuple orderings, coordinate permutations, carrier labels, and

equivalent embeddings. Without reduction to canonical representatives, raw counts would measure presentations rather than structures.

The canonical form supplies a stable identifier for:

deduplicating candidate lifts,

matching outputs between independent implementations,

organizing atlas rows,

checking determinism,

and verifying that repeated runs return the same representatives.

The use of exhaustive permutation-based canonicalization is computationally practical at orders two, three, and four. It would become a serious bottleneck at larger carrier sizes, where more sophisticated graph-canonicalization or partition-refinement methods would be required.

At the present scale, exhaustive canonicalization has the advantage of transparency. No heuristic choices are needed, and every equivalence under relabeling is checked directly.

5.3 Exact Two-Element Variety Membership

The sixteen labeled operations and ten isomorphism types do not automatically correspond to ten distinct equational theories. Nonisomorphic finite algebras can satisfy exactly the same identities. The program must therefore determine the directed variety-membership relation among the two-element operations.

For two cores B and C, the question is whether

$B \in V(C)$.

Equivalently, every identity true in C must also be true in B.

In general, deciding variety membership can be difficult. At the present two-element scale, the relation can be computed exactly through finite powers and term functions.

Every term in the binary signature induces a function on a two-element algebra. For k variables, a term operation is a map

$B^k \to B$.

Although infinitely many syntactic terms exist, only finitely many such functions exist. The term clone is obtained by beginning with projection functions and closing under pointwise application of the basic binary operation.

For two-element algebras, a finite arity bound suffices for the exact membership computation used in the artifact. The implementation generates the relevant term operations, constructs the induced free-algebra behavior, and determines whether one algebra belongs to the variety generated by another.

The result is a complete directed 16 × 16 membership matrix.

This matrix is then used to compute equivalence classes under mutual variety membership. Two operations B and C belong to the same exact law class when

$B \in V(C)$

and

$C \in V(B)$.

Equivalently,

$V(B) = V(C)$,

and therefore

$Eq(B) = Eq(C)$.

The computation finds ten exact law classes among the sixteen labeled operations.

At this order, the ten law classes coincide with the ten isomorphism classes. This is a result of the computation, not an assumption. It says that no two nonisomorphic two-element binary algebras in the catalog generate the same variety.

The coincidence simplifies the four-element atlas because each isomorphism type can be treated as a distinct equational core. At larger orders, such coincidence should not be expected.

5.4 Enumeration of Three- and Four-Element Residual Lifts

The primary lift enumeration proceeds by examining coordinate realizations inside powers of each two-element core.

For a selected core B and a coordinate dimension m, the ambient algebra is B^m with componentwise operation. A candidate lift is a subset

$S \subseteq B^m$

satisfying three conditions.

First, S has the desired carrier size, either three or four.

Second, S is closed under the componentwise operation:

$x,y \in S$ implies $x \star y \in S$.

Third, each coordinate projection maps S onto all of B.

Closure makes S a subalgebra. Surjective coordinate projections make the embedding subdirect.

For a four-element lift, the pair-separation principle gives a universal finite coordinate bound. There are six unordered pairs of distinct elements in a four-element set. If the algebra is residually B, then each pair can be separated by some homomorphism to B. Selecting one coordinate for each pair yields an embedding into B^6.

Thus every residual four-element lift appears in some power B^m with

$m \leq 6$.

The program can therefore search finite coordinate dimensions without assuming that two or three coordinates will always suffice.

In practice, the classification focuses on minimal separating families and discovers that the surviving four-element types require either two or three coordinates. No classified case requires four, five, or six minimum views.

Three-element lifts are also enumerated as an internal structural layer. They provide additional checks on the construction, help verify the residual catalogs, and make the search architecture less dependent on one target cardinality.

The enumerator examines coordinate families, constructs their images as subalgebras, canonicalizes the induced operation tables, and records the relationship between each lift and its core.

The released classification contains seven nonisomorphic three-element residual core-lift pairs and sixteen nonisomorphic four-element residual core-lift pairs.

5.5 Canonicalization Under Relabeling

A single abstract lift can be presented in many ways.

Its four elements can be assigned any of 24 labelings.

Its coordinate positions can be permuted.

The values 0 and 1 may be relabeled within a core when an automorphism permits it.

Different faithful families may embed the same algebra into different powers or into different subsets of the same power.

The computation must identify which presentations represent the same algebraic object.

Each induced lift operation is therefore converted into a canonical operation table under all carrier permutations. The canonical table becomes the primary isomorphism key.

Additional information, such as the selected core type and quotient-family geometry, is then attached to the canonical lift.

Coordinate permutations do not need to be preserved as distinct structures unless they induce genuinely different quotient-family data. A reconstruction using the first and third coordinates is not intrinsically different from one using the second and first if an automorphism or coordinate relabeling carries one to the other.

Canonicalization also allows the primary and audit implementations to be compared field by field. The two paths need not generate candidates in the same order or use the same temporary labels. They need only reduce to the same canonical structures and invariants.

Deterministic serialization is essential. Lists, partitions, policies, and operation tables are sorted by explicit rules before being written. This permits byte-for-byte reproduction rather than merely numerical agreement.

At small order, the canonicalization is exhaustive enough to serve as a direct computational proof of nonisomorphism among distinct atlas rows.

5.6 Verification of Quotient Maps and Kernel Intersections

For each candidate lift A and core B, the program enumerates surjective homomorphisms

$q : A \to B$.

Since A has four elements and B has two, there are only

$2^4 = 16$

set maps from A to B, of which 14 are surjective. Each can be tested directly against the homomorphism condition.

Every surviving homomorphism determines a two-block kernel partition. Equivalent maps differing only by an automorphism of the core may induce the same kernel. The program records both map-level and policy-level information as needed.

For each kernel θ, the induced quotient algebra A/θ is computed and checked for isomorphism with B. This verifies explicitly that the map is not merely surjective but yields the intended core type.

The admissible quotient kernels are then combined in families. For a family

$\theta_1,\ldots,\theta_m$,

the program computes their intersection. The family separates A exactly when

$\theta_1 \cap \ldots \cap \theta_m = \Delta A$.

The minimum family size is found by checking subsets in increasing cardinality.

For four-element lifts, no one-map family can be faithful because every map onto a two-element core identifies at least two pairs of elements. The first possible minimum is therefore two.

If no pair of kernels intersects diagonally but some triple does, then the core-view number is three.

This exact kernel calculation produces the 10/6 division:

ten lift types admit a faithful pair,

six require a faithful triple.

The direct-square versus proper-cube distinction follows immediately from the minimum family size.

5.7 Minimum Reconstruction Families

Finding the minimum number of views is only the first step. The program also identifies every minimum family.

Suppose A has r admissible core quotient policies. The candidate minimum families are subsets of the r policies. A subset is accepted when its kernel intersection is diagonal and no proper subfamily has diagonal intersection.

For each accepted family, the computation records:

its participating quotient policies,

its cardinality,

the tuple representation of A,

the induced coordinate geometry,

and its orbit under Aut(A).

This information defines the reconstruction hypergraph.

The enumeration of all minimum families allows several stronger conclusions than the view number alone.

It determines whether reconstruction is unique or redundant.

It reveals whether certain quotient policies are indispensable.

It shows how different minimum families overlap.

It allows automorphism-orbit analysis.

It distinguishes arbitrary coordinate choices from invariant reconstruction structure.

In all sixteen four-element cases, the minimum families form a single orbit under the automorphism group of the lift. This means that although several labeled minimum families may exist, they are structurally equivalent.

The program therefore avoids selecting one family as uniquely fundamental. It reports the orbit structure and, where useful, a canonical representative for serialization.

5.8 Independent Audit and Reproducibility Architecture

The artifact includes a second computational path that reimplements the central classification logic separately from the primary enumerator.

Its purpose is not to provide socially independent verification by an external researcher. It is an internal independence check at the level of implementation path.

The audit recomputes:

the complete two-element variety-membership matrix,

the partition into isomorphism classes,

the partition into exact law classes,

the three-element residual lift catalog,

the four-element residual lift catalog,

the core/lift/minimum-view triples,

the released atlas fields,

and the independent four-element meet-semilattice enumeration.

The audit then compares its results with the frozen release outputs and writes an explicit report.

This design reduces the risk that one implementation error could generate both the data and the apparent verification. The two paths may still share conceptual assumptions, definitions,

and parts of the surrounding environment, so the audit is not infallible. But agreement between separately organized algorithms is much stronger than rerunning one function twice.

Reproducibility is reinforced by additional layers.

The release reproduction command regenerates the principal data files into a designated output directory and compares them byte-for-byte with the frozen files.

The audit requires an explicit results directory so that an installed package does not silently assume that repository fixtures are present.

The integrity verifier checks file hashes, byte sizes, manifest agreement, and exact inclusion of release files.

The determinism checker confirms that repeated generation produces identical outputs.

The test suite verifies normal execution, optimized Python execution, corruption rejection, manifest rejection of unexpected files, environment sanitization, canonicalization, and packaging behavior.

The verification environment pins exact known-good toolchain versions while the ordinary package metadata allows a wider compatible installation range.

Together, these layers distinguish four questions:

Can the code run?

Can it reproduce the frozen findings?

Can an independent implementation path recover the same classification?

Can alteration of the frozen artifact be detected?

A strong computational release should answer all four.

5.9 Scope, Completeness, and Computational Boundaries

The computation is exhaustive only within its declared class.

It covers:

finite algebras with one binary operation,

two-element cores,

three- and four-element residual lifts,

surjective homomorphisms onto a fixed core,

subdirect reconstruction through core-valued quotient views,

exact equality of equational theories established structurally,

and classification up to isomorphism.

It does not cover:

all four-element binary algebras,

all four-element algebras equationally equivalent to a two-element quotient,

nonresidual exact extensions,

mixed-core reconstruction systems,

multiple-operation signatures,

relations without operations,

larger carriers,

quasi-equational or full first-order equivalence,

approximate homomorphisms,

or infinite structures.

The phrase exhaustive classification must therefore always be followed, explicitly or implicitly, by its scope:

the four-element binary algebras residually reconstructible from a two-element exact law core.

Within that domain, the classification is complete. Outside it, the present results make no census claim.

The computational method also faces scaling limits. Exhaustive carrier permutation grows factorially. Candidate coordinate families grow combinatorially. The number of operation tables grows as $n^{(n^2)}$. Homomorphism, congruence, and subalgebra counts can increase rapidly.

The order-four experiment succeeds because the search is core-first, the coordinate bound is finite, and all relevant structures remain small enough for exact enumeration.

Future extensions will require stronger symmetry reduction, canonical augmentation, congruence-guided generation, parallel search, checkpointing, and perhaps formal proof extraction.

The present computation should therefore be regarded as a complete base case rather than a universal solver. Its importance lies in demonstrating that the theory-preserving ontogenesis question can be made exact, exhaustively computed, independently audited, and expressed through invariants that reveal structure beyond raw cardinality.

6. The Exact Four-State Classification

The completed classification identifies every four-element binary algebra in the declared residual universe that can be reconstructed from a two-element exact law core. The result is finite, explicit, and structurally sharper than a list of examples. It organizes the surviving algebras by core type, minimum number of lawful quotient views, quotient geometry, automorphism behavior, splitting structure, and familiar algebraic properties.

The central count is:

10 nonisomorphic two-element cores,

16 nonisomorphic four-element residual lifts,

10 lifts with core-view number two,

6 lifts with core-view number three.

The distinction between the last two groups is the principal structural result. Every four-element set can be injected into a two-coordinate binary space. Yet six of the classified algebras cannot be reconstructed using any pair of lawful core quotient maps. They require three.

This division is not imposed by the search design. The computation allows larger coordinate families and tests every admissible quotient combination. The 10/6 split emerges from the kernel geometry of the classified lifts.

The atlas also shows that several features vary independently. Two lifts over the same core may differ in reconstruction number, automorphism group, quotient count, or operation properties. Conversely, lifts with different familiar descriptions may share the same minimum reconstruction geometry. The classification is therefore not reducible to a catalog of named operations.

It is an atlas of how one complete equational theory can be embodied in a larger finite world.

6.1 The Ten Two-Element Equational Cores

The starting universe contains sixteen labeled binary operations on a two-element carrier. Relabeling the two carrier elements identifies some of these operations, reducing the catalog to ten isomorphism types.

A separate exact computation of variety generation also produces ten law classes. In this smallest universe, isomorphism type and complete equational-theory type coincide.

This coincidence is useful but should not be generalized without proof. Nonisomorphic finite algebras can generate the same variety at larger cardinalities. Here, however, each of the ten two-element isomorphism types supplies a distinct exact law core.

The cores include several recognizable behaviors:

constant operations,

left and right projections,

semilattice-like operations,

group-like or parity-like operations,

and asymmetric operations lacking the standard combinations of associativity, commutativity, and idempotence.

The classification treats all ten uniformly. A core does not qualify because it resembles a familiar algebraic textbook object. It qualifies because it is one of the complete nonisomorphic two-element binary structures.

Each core determines its own universe of residual lifts. Some cores are more fecund than others: they support multiple nonisomorphic four-element realizations. Others support few or none within the specified residual and quotient constraints.

This variation is itself meaningful. Equal carrier size does not imply equal generative capacity. Two cores may both contain two elements yet differ substantially in the number and geometry of larger theory-preserving worlds they admit.

The fecundity summary records this distribution. It is a first indication that the law carried by the operation, rather than the cardinality of the carrier alone, governs ontogenetic possibility.

6.2 The Sixteen Nonisomorphic Four-Element Lift Types

The exhaustive residual search produces sixteen nonisomorphic four-element core-lift types.

A type records more than a four-element operation table. It includes the relation between the lift and its designated two-element core, the admissible quotient policies, the minimum faithful families, the reconstruction dimension, and the relevant symmetry data.

Two four-element algebras with isomorphic operation tables but different core assignments would not automatically represent the same core-lift relation. In the present atlas, the classification is organized so that the chosen core is part of the structure being reported.

Each lift satisfies three essential conditions:

it maps surjectively onto its selected core,

it embeds subdirectly into a finite power of that core,

and it shares exactly the same equational theory with the core.

The sixteen types therefore represent sixteen distinct ways in which a two-element law world can be realized on four elements without changing its complete theory.

The number sixteen should not be confused with the sixteen labeled two-element operations. The equality is numerical coincidence. One count concerns all labeled core operations before isomorphism reduction; the other concerns nonisomorphic four-element residual lift types after full classification.

The atlas is small enough for complete inspection but large enough to show that theory-preserving enlargement is not unique. A given equational theory can support multiple larger embodiments with different internal structures.

This is one of the clearest mathematical senses in which ontogenesis occurs. New objects and new relational organization appear without the birth of a new equational theory.

6.3 The Ten Direct-Square Reconstructions

Ten of the sixteen four-element lifts have core-view number two.

For such a lift A over a two-element core B, there exist two surjective homomorphisms

$q_1, q_2 : A \to B$

whose kernels intersect diagonally.

The combined map

$a \mapsto (q_1(a), q_2(a))$

embeds A into B^2.

Since both A and B^2 contain four elements, the injective map is also surjective. Therefore

$A \cong B^2$.

Every two-view four-element lift is thus the direct square of its core.

This is not merely a convenient representation. It is forced by cardinality. There is no room for a proper four-element subalgebra inside a four-element ambient square.

The ten direct-square cases therefore have a particularly transparent structure. Their elements correspond to all four possible coordinate pairs:

(0,0), (0,1), (1,0), (1,1).

The operation is performed coordinatewise according to the core operation.

Yet even in these cases, the quotient geometry may contain more than the two obvious coordinate projections. Additional core quotients can exist, and multiple minimum pairs may reconstruct the same lift. The reconstruction hypergraph records these alternatives.

The square cases establish the baseline form of theory-preserving enlargement. The core is duplicated into two coordinated views, and the full product realizes a four-element world with the same equational theory.

In information-theoretic terms, these cases achieve the cardinal minimum exactly. Two lawful binary views are both necessary and sufficient.

6.4 The Six Proper-Cube Reconstructions

The remaining six lift types have core-view number three.

For each such lift A over B, no pair of surjective homomorphisms onto B has diagonal kernel intersection. Every lawful pair leaves at least one pair of elements indistinguishable.

A faithful triple exists, however. The diagonal map embeds A into B^3:

$A \leq B^3$.

Because B^3 has eight elements while A has four, the image is a proper subalgebra rather than the full cube.

These six cases are therefore proper-cube reconstructions. Their four elements occupy only half of the available binary triples, but the selected subset is closed under the coordinatewise core operation and projects surjectively onto every coordinate.

The geometry is more constrained than arbitrary four-point selection. Most four-point subsets of an eight-point cube are not subalgebras. Others may be subalgebras but fail one or more

coordinate projections. Others may project surjectively yet admit a redundant coordinate, reducing the true view number to two.

The six classified cases survive all these tests. Their three coordinates are jointly necessary at minimum.

The proper-cube cases demonstrate that carrier size does not determine reconstruction dimension. A four-element algebra may occupy a three-dimensional lawful coordinate system even though two dimensions provide enough raw labels.

The missing cube points matter. They reflect combinations of core states that are algebraically unavailable in the lift. The lift is not the unconstrained product of three core copies. It is a closed relational pattern within that product.

This restricted pattern is precisely what forces the third view. The available quotient maps cannot be compressed into two lawful coordinates without merging distinct elements.

6.5 Splitting of Exact Core Quotient Maps

A surjective homomorphism

$q : A \to B$

splits when there exists a homomorphism

$s : B \to A$

such that

$q \circ s = idB$.

The map s is a section of q. It embeds a copy of the core back into the lift and selects one representative from each quotient class in a way compatible with the operation.

Splitting is stronger than the mere existence of a quotient. A quotient can exist without any homomorphic copy of the quotient algebra being embedded as a section inside the source.

Within all sixteen classified four-element cases, every exact core quotient map splits.

This is a striking regularity of the atlas. Every lawful collapse onto the core admits a lawful return embedding of that core into the larger world.

The result means that the core is not only recoverable as a quotient abstraction. It is also present internally as a subalgebra selected by a section.

The lift therefore contains the core in two complementary senses:

externally, as a quotient produced by identification;

internally, as an embedded subalgebra returned by a splitting.

This does not make the lift a trivial product in every case. The six proper-cube structures remain non-square reconstructions despite the splitting of each quotient. Splitting one quotient does not imply that two quotients suffice jointly.

The universal splitting found here may be special to the tiny residual universe. It should not be assumed for larger carriers or broader classes. Future searches should test whether nonsplit exact core quotients appear and whether they introduce new forms of ontogenetic asymmetry.

6.6 Automorphism Orbits of Minimum Reconstruction Families

A lift may possess several minimum faithful quotient families. The raw number of such families depends partly on labeling, so their intrinsic meaning must be evaluated under the automorphism group of the lift.

For each of the sixteen types, the computation determines the action of Aut(A) on the set of minimum reconstruction families.

In every case, all minimum families lie in a single automorphism orbit.

This means that any minimum reconstruction can be transformed into any other by a symmetry of the lift. No minimum family is structurally preferred.

The result is important conceptually. A classification based on one chosen coordinate representation could falsely suggest that a particular set of views is fundamental. Orbit analysis shows instead that the lift determines an equivalence class of minimum reconstructions.

The algebra has an intrinsic reconstruction geometry but not necessarily an intrinsic coordinate frame.

This parallels a general principle in mathematics and physics: coordinate descriptions may vary while invariant structure remains fixed. The appropriate object is the orbit of descriptions under symmetry.

In the present setting, the invariant may be expressed as the reconstruction hypergraph equipped with the action of Aut(A). A listed minimum family is merely one representative.

The single-orbit property also suggests a form of internal coherence. The lift does not possess several disconnected kinds of minimum reconstruction. Its lawful minimal perspectives form one symmetry-connected class.

Whether this remains true beyond four elements is an open computational question. Larger systems may exhibit multiple inequivalent reconstruction orbits, creating a richer notion of perspectival phase or reconstruction regime.

6.7 Independent Recovery of the Five Meet-Semilattices

A meet-semilattice is an algebra with a binary operation satisfying:

associativity,

commutativity,

and idempotence.

These three identities give the operation an order-theoretic interpretation as meet.

The four-element classification contains exactly five meet-semilattice types. Their recovery serves as an important internal validation because they can be obtained in two conceptually different ways.

The general residual-lift enumeration discovers them without assuming semilattice laws in advance. It begins from anonymous binary cores and lawful coordinate constructions.

A separate semilattice-specific enumeration generates four-element meet-semilattices directly from their defining identities or corresponding order structures.

The two paths produce exactly the same five isomorphism types.

This agreement tests more than a headline count. It shows that the general ontogenesis pipeline correctly recovers a familiar, independently characterizable subclass embedded within the anonymous-operation universe.

The semilattice cases also illuminate the larger philosophical program. Familiar mathematical objects can appear after the computation rather than being inserted at the beginning. The search does not start by asking for five known semilattices. It begins with unnamed finite operations and later recognizes that a subset satisfies the semilattice laws.

This is a modest but genuine instance of theory recognition following structural emergence.

6.8 The Structure of the Complete Classification Atlas

The released atlas records one row for each of the sixteen four-element core-lift types. Its fields combine algebraic description, reconstruction data, and reproducibility identifiers.

The recorded information includes:

the canonical identity of the two-element core,

the canonical identity of the four-element lift,

the lift operation table,

the minimum core-view number,

the number of admissible core quotient policies,

the minimum reconstruction families,

the automorphism group size,

the number and orbit structure of minimum families,

the splitting behavior of quotient maps,

the direct-square or proper-cube status,

and selected familiar operation properties.

The atlas is supplemented by several additional tables.

The core fecundity summary groups lift counts by core.

The exact law-class table records the partition of the sixteen labeled two-element operations.

The variety-membership matrix records the complete directed relation among the two-element operations.

The machine-readable summary consolidates the principal counts and verification status.

Together these files provide more than a narrative statement. They expose the complete finite evidence supporting the classification.

The atlas also makes future comparison possible. A larger-order computation can preserve the same field logic while extending the carrier size, operation signature, or permitted core families. New invariants can be added without losing continuity with the base case.

The four-state atlas should therefore be regarded as both a finished classification and a reference coordinate system for subsequent work.

Its most consequential division remains the simplest:

ten square worlds,

six cube worlds.

The next section examines why the six proper-cube cases matter and what they reveal about the difference between raw information and lawful mathematical reconstruction.

7. The Third-View Phenomenon

The most important mathematical result in the four-state classification is the discovery that six lawful object-worlds require three binary quotient views for exact reconstruction, even though four objects can always be distinguished by two unrestricted binary coordinates. This gap separates information capacity from algebraic admissibility.

At the level of cardinality, the problem appears trivial. A set with four elements can be injected into the binary square:

(0,0), (0,1), (1,0), (1,1).

Two binary coordinates therefore carry enough raw information to assign a unique label to every element. If reconstruction meant only that each object received a different bit string, no third coordinate could ever be necessary.

Theory-Preserving Ontogenesis imposes a stronger requirement. Every coordinate must be a surjective homomorphism onto the same two-element law core. Each coordinate must preserve operation, not merely distinguish symbols. Consequently, most possible binary labelings are inadmissible.

The six proper-cube cases show that lawful coordinates can be too constrained to realize the cardinal minimum. No pair of admissible quotient views distinguishes every pair of elements. A triple does.

The third view is therefore not added because the system has too many objects. It is added because the available lawful distinctions overlap in a way that prevents any two from being jointly faithful.

This phenomenon provides the sharpest finite evidence for the claim that objecthood depends on coherent relational structure rather than labels alone.

7.1 The Two-Bit Cardinality Bound

For finite sets X and Y, an injective encoding

$X \to Y^m$

requires

$|Y|^m \geq |X|$.

When $|X| = 4$ and $|Y| = 2$, the minimum unrestricted coordinate count is two because

$2^1 < 4 \leq 2^2$.

This is a purely cardinal statement. It does not refer to operations, laws, congruences, or homomorphisms. It says only that four distinct codewords are available in a two-coordinate binary space.

The bound is exact for arbitrary set-theoretic encoding. Every four-element set can be placed in bijection with the full binary square.

The direct-square lifts achieve this bound lawfully. Their four elements correspond to all four binary pairs, and both coordinate projections are surjective homomorphisms onto the core.

The proper-cube lifts do not. Their minimum lawful embedding has the form

$A \to B^3$,

with $|A| = 4$ and $|B^3| = 8$.

The image uses only half of the available codewords. Thus the extra coordinate does not increase the number of represented objects. It changes which lawful distinctions can coexist.

One can express the contrast through two quantities.

The unrestricted coordinate number is

$\lambda B(A) = \text{ceiling}(\log|B| \, |A|)$.

The lawful coordinate number is the core-view number

$\nu B(A)$.

For the classified four-state systems,

$\lambda B(A) = 2$

in every case, while

$\nu B(A) \in \{2,3\}$.

The difference

$\sigma B(A) = \nu B(A) - \lambda B(A)$

is zero for ten lifts and one for six lifts.

This difference may be called the lawful reconstruction overhead. It measures the coordinate cost imposed by preservation of algebraic structure beyond the minimum required by cardinality.

7.2 Unrestricted Labeling Versus Lawful Observation

An unrestricted binary coordinate is any function

$f : A \to \{0,1\}$.

For a four-element set, there are sixteen such functions. Many pairs of them combine into injective encodings.

A lawful binary coordinate is much rarer. It must satisfy

$f(a \star_A b) = f(a) \star_B f(b)$

for all a,b in A.

It must also be surjective if it is to represent a complete quotient view of the core.

These equations constrain the values of f across the entire operation table. Choosing the image of one element can force the images of others. A partition that is suitable for labeling may fail compatibility with the operation. The quotient classes may not be closed under the induced behavior required by B.

Thus the set of admissible views is not the set of all binary questions one could ask about A. It is the set of binary questions whose answers transform coherently under the operation.

This distinction can be phrased conceptually:

An unrestricted coordinate reports a distinction.

A lawful coordinate participates in the algebra.

The second is stronger because it must make the following diagram commute for every pair of inputs:

combine in A, then observe;

or observe each input, then combine in B.

Both routes must yield the same core value.

A lawful view therefore preserves not only states but transitions of combination. It is an operational perspective rather than a static label.

This is why ordinary coding theory supplies only a lower bound. The available codewords may be numerous enough, but the permitted coordinate functions may be too structured to realize them in the required combinations.

7.3 Why Two Binary Views Can Fail

Suppose A has several surjective homomorphisms onto B. Each homomorphism has a kernel partition with two blocks. Because A has four elements, every such partition identifies at least one pair of distinct elements.

Two quotient maps q_1 and q_2 reconstruct A exactly when no pair remains identified by both. Equivalently,

$\ker(q_1) \cap \ker(q_2) = \Delta A$.

Failure occurs when every pair of available quotient kernels has a nontrivial intersection. That is, for every choice of two lawful views, at least one distinct pair of elements is collapsed by both.

The failure does not imply that the quotient maps are poor individually. Each is a valid surjective homomorphism. Each presents the core theory correctly. The problem is coordination.

The available kernel partitions may overlap in a structured pattern. One view distinguishes some pairs but collapses others. A second view resolves many of those collapses but necessarily repeats at least one of them. Only a third view breaks the remaining ambiguity.

This is a covering problem over pair distinctions.

Let P(A) denote the six unordered pairs of distinct elements of a four-element carrier. Each quotient map q separates a subset

$S(q) \subseteq P(A)$,

consisting of pairs whose images differ.

A faithful family must satisfy

$S(q_1) \cup \ldots \cup S(q_m) = P(A)$.

Two-view failure means that no two admissible separation sets cover all six pairs. Three-view success means that some triple does.

The third-view phenomenon can therefore be studied either through intersections of kernels or unions of separation sets. The two descriptions are dual.

The kernel language emphasizes identifications.

The separation language emphasizes distinctions.

Both reveal that the obstruction is combinatorial but generated by algebraic law.

7.4 Kernel Geometry and Incomplete Separation

The set of congruences of an algebra forms a lattice under intersection and generated join. Core quotient policies occupy a selected region of this congruence lattice: those congruences whose quotient is isomorphic to the designated core.

The core-view number is governed by how these congruences intersect.

If two core congruences meet at the diagonal, the lift is a square reconstruction.

If every pair meets above the diagonal but some triple meets at the diagonal, the lift is a cube reconstruction.

Thus the difference between two-view and three-view worlds is visible in the meet geometry of their core congruences.

The three-view cases have no complementary pair among the admissible core kernels. The quotient lattice may contain other congruences, but the restricted family of exact core quotients lacks a diagonally intersecting pair.

This restriction is essential. A pair of arbitrary partitions might distinguish all elements. A pair of congruences of other quotient types might also do so. The invariant asks specifically for quotients onto the same exact law core.

The geometry is therefore policy-indexed. Reconstruction depends not only on the congruence lattice of A but on which congruences qualify under the selected core.

The third view arises from the failure of pairwise complementarity within that admissible family.

This suggests a general hierarchy. Larger algebras may require four or more lawful views when every smaller subfamily retains a nontrivial common kernel. The core-view number then measures the depth at which the meet of admissible core congruences reaches equality.

7.5 Six Worlds Requiring Three Theory-Preserving Views

The exhaustive classification finds exactly six four-element residual lift types with

$vB(A) = 3$.

For each of these lifts:

at least three surjective homomorphisms onto B exist;

some triple has diagonal kernel intersection;

no pair has diagonal kernel intersection;

the diagonal image is a four-element proper subalgebra of B^3;

and A satisfies exactly the same identities as B.

The number six is not obtained by selecting six illustrative examples. It is the complete count within the classified residual universe.

These six worlds demonstrate several points simultaneously.

First, lawful reconstruction overhead exists at the smallest scale where it can be cleanly observed.

Second, the overhead is not exceptional in the sense of occurring only once. It appears in more than one-third of the sixteen classified lift types.

Third, the effect survives exact theory preservation. The third coordinate is not required because the lift has acquired a different or more complicated equational theory. Core and lift remain equationally equivalent.

Fourth, the proper-cube form is not reducible to a poorly chosen presentation. The failure of every lawful pair is invariant under isomorphism and automorphism.

The six cases therefore constitute a genuine structural class, not an artifact of labeling.

Their existence is the central reason the project goes beyond a routine enumeration of small algebras. They expose a measurable difference between two notions that ordinary language often treats as equivalent:

enough information to identify an object;

enough lawful perspectives to reconstruct it.

7.6 Algebraic Structure as the Source of the Extra View

Why is the third coordinate forced? The answer is not that binary values are scarce. Eight triples are available, and only four are used. Nor is it that the operation requires more output symbols. Every coordinate uses the same two-element core.

The source is the compatibility condition linking the coordinate partitions to the operation.

Each lawful quotient kernel must be a congruence. This means that equivalence classes cannot be chosen freely. They must remain stable under pairwise combination.

The operation therefore shapes the possible observational cuts through the carrier. It decides which distinctions can be made together in one homomorphic view.

Some operations admit two congruences whose overlap is only equality. Their lifts are direct squares.

Other operations restrict admissible congruences so that every pair overlaps nontrivially. Their lifts require a third quotient.

The extra view is consequently a property of the relational organization encoded by the operation table.

This leads to a general principle:

Coordinate cost is determined not only by how many states exist, but by which state distinctions are compatible with law.

The principle resembles constraints in error-correcting codes, distributed sensing, tomography, and contextual measurement, but the present result is exact within finite universal algebra. The permissible coordinates are homomorphisms, and the obstruction is computed through congruence intersections.

7.7 Being Distinguishable and Being Reconstructible

A set of labels can distinguish objects without explaining how they behave. Two elements with different labels may have no lawful relationship between their labels and the operation of the system.

Reconstruction requires more. It must recover both identity and operation.

Given an injective arbitrary code

$c : A \to \{0,1\}^2$,

one could transport the operation of A to the four codewords. But the transported operation need not act coordinatewise according to the selected two-element core. The bits may interact in a coupled way that destroys their interpretation as independent lawful views.

By contrast, a homomorphic diagonal representation satisfies

$c(a \star_A b) = c(a) \star c(b)$

coordinate by coordinate.

The operation of the lift is then recoverable entirely from the core operation applied separately in each view.

Thus a lawful reconstruction is compositional. Each coordinate has a stable local meaning, and the global operation is assembled from those meanings.

The third-view cases show that two arbitrary labels can encode the carrier but cannot decompose its operation into two admissible core perspectives.

This yields the distinction:

Distinguishability concerns uniqueness of code.

Reconstructibility concerns uniqueness together with lawful compositional recovery.

The second notion is closer to mathematical objecthood because it preserves the relations that make the object part of an algebraic world.

7.8 Information Capacity Constrained by Law

Information theory ordinarily studies what can be represented or transmitted given an alphabet and channel. Theory-Preserving Ontogenesis adds a law constraint to the alphabet.

The core symbols are not neutral. They combine according to a fixed operation. A coordinate is valid only when observation commutes with combination. The channel is therefore algebraically structured.

One may define a lawful coding problem:

Given algebras A and B, determine the minimum m for which A embeds subdirectly into B^m.

The ordinary cardinal lower bound remains necessary:

$|B|^m \geq |A|$.

But it is no longer sufficient.

The ratio between lawful and unrestricted coordinate requirements may become an informative invariant at larger scales. One can ask whether the overhead remains bounded, grows logarithmically, grows linearly, or depends sharply on the identities of the core.

The four-state result provides the first nonzero example in the present program:

2 unrestricted views,

3 lawful views.

The gap is only one coordinate, but conceptually it is decisive. It proves that law can impose representational costs invisible to cardinality.

The result also suggests that mathematical information should sometimes be measured relative to admissible morphisms rather than arbitrary functions. A system may contain enough nominal capacity while lacking enough lawful channels.

7.9 The Third View as Relational Closure

The third coordinate can be interpreted not merely as additional information but as relational closure.

Two views may nearly reconstruct the lift. They may distinguish five of the six element pairs while failing on one. The third view closes the remaining gap.

Yet its role is not necessarily confined to one pair. Because the minimum families form symmetry orbits, each coordinate may participate equally in the total geometry. Removing any one may reintroduce ambiguity.

The triple then functions as an irreducible relational whole. No coordinate alone is complete, and no pair is sufficient. Identity emerges only at the level of the full family.

This structure resists a simple hierarchy in which one primary view is supplemented by two secondary corrections. The automorphism action may show that all three are equivalent under symmetry.

The object-world is therefore reconstructed through a triadic relation among perspectives.

This should not be turned prematurely into a theological numerology. The number three in these finite algebras does not prove or encode the Trinity. Its significance is mathematical: minimum faithful reconstruction can be irreducibly triadic.

The philosophical lesson is more general. Complete objecthood may arise only from a family of mutually limited but jointly sufficient views.

7.10 Why the Phenomenon Matters for Ontogenesis

The third-view phenomenon converts the language of lawful emergence into a concrete mathematical difference.

If every four-state lift had required exactly two views, one might argue that the entire theory-preserving framework merely repackaged binary coding. The quotient maps would add terminology without changing the reconstruction count.

The six cube cases refute that reduction. Algebraic law changes the answer.

The larger object-world cannot always be recovered by the smallest number of labels. Its reconstruction requires the smallest number of lawful relations.

That is the core ontogenetic finding:

A mathematical object-world is not determined solely by how many elements it contains, but by how its admissible perspectives intersect.

The result does not establish a universal metaphysics of objecthood. It establishes a precise finite instance in which relation and law contribute irreducibly to reconstruction.

The next section develops the invariants that capture this contribution: the core-view number and the reconstruction hypergraph.

8. Core-View Number and Reconstruction Hypergraph

The third-view phenomenon shows that lawful reconstruction requires invariants more refined than carrier size. Cardinality determines how many unrestricted symbols are necessary to distinguish elements, but it does not determine how many quotient views are required when every view must preserve the operation and land in the same exact law core.

Two complementary invariants capture this richer structure.

The first is the core-view number. It records the minimum number of lawful core-valued quotient maps whose kernels intersect only in equality.

The second is the reconstruction hypergraph. It records the entire family of admissible quotient policies and identifies which minimum subsets are jointly faithful.

The core-view number compresses the reconstruction problem into one integer. The reconstruction hypergraph retains the geometry hidden by that compression.

Together they answer two different questions:

How many lawful views are necessary?

Which lawful views can cooperate to achieve reconstruction?

This distinction parallels a common pattern in mathematics. A dimension records the size of a minimum basis, but the geometry of all bases contains additional structure. A rank records the minimum number of generators, but the dependency relations among possible generators form a matroid or related combinatorial object. A covering number gives a minimum count, but the family of all minimum covers reveals redundancy and symmetry.

The reconstruction problem has the same layered character. The number alone is important, but the organization of all minimum faithful families may be the deeper invariant.

8.1 Reconstruction Number as a Mathematical Invariant

For a lift A and a selected exact core B, the core-view number is

vB(A)

min {m : there exist surjective homomorphisms $q_1,\ldots,q_m : A \to B$ with $\bigcap_i \ker(q_i) = \Delta A$}.

This quantity is invariant under isomorphism of the pair.

Suppose

$f : A \to A'$

and

$g : B \to B'$

are algebra isomorphisms. Any quotient map

$q : A \to B$

induces a corresponding quotient map

$g \circ q \circ f^{-1} : A' \to B'$.

Kernel intersections are preserved under transport by f. Therefore

$vB(A) = vB'(A')$.

The number is thus independent of carrier labeling, coordinate order, or the chosen canonical presentation.

It measures a relative property. The notation includes B because reconstruction is indexed by the selected core. An algebra A might admit several exact cores, and the minimum number of views could differ among them.

This relativity is not a weakness. It reflects the fact that resolution depends on the language of observation. The same larger world may be easy to reconstruct through one core and difficult through another.

The invariant also separates two dimensions of complexity.

Carrier complexity is reflected by $|A|$.

View complexity is reflected by $vB(A)$.

These quantities can diverge. Two algebras of equal cardinality may require different numbers of lawful views. Conversely, algebras of different sizes might share the same view number relative to different cores.

The four-state classification gives the first complete distribution:

$vB(A) = 2$ for ten lift types,

vB(A) = 3 for six lift types.

This distribution converts a qualitative observation into a finite invariant census.

8.2 Relation to Minimal Subdirect Decomposition

The core-view number belongs to the broader mathematical territory of subdirect representations. Universal algebra has long studied embeddings of an algebra into products of quotient algebras and decompositions into subdirectly irreducible factors.

The present invariant is more restricted than an arbitrary minimum subdirect decomposition.

It does not permit any convenient family of factors.

It fixes one core B.

Every factor must be isomorphic to B.

Every coordinate must be surjective.

The quotient must preserve the exact equational theory.

The objective is not merely to decompose A but to reconstruct it from repeated views of one law-bearing world.

This restriction is what gives the invariant its ontogenetic interpretation. The question is not how economically A can be represented using arbitrary simpler algebras. It is how many copies of its own exact core theory are required to recover it.

In ordinary subdirect decomposition, different factors may capture different portions of the algebra's law. One factor might be group-like, another semilattice-like, and another trivial. In fixed-core reconstruction, every coordinate speaks the same algebraic language.

The resulting minimum can therefore exceed the minimum number of arbitrary quotient factors.

This provides a useful calibration of novelty. The broad idea of minimizing product factors is classical. The distinctive contribution lies in fixing an equationally exact core and treating the resulting minimum as an organizing invariant of theory-preserving enlargement.

The core-view number can be regarded as a constrained subdirect dimension.

8.3 The Importance of Fixing an Exact Law Core

Why insist that every quotient be isomorphic to one fixed core?

Without this restriction, reconstruction could combine incomparable perspectives. A family of quotients might separate all elements efficiently while distributing the source theory across different factor types. Such a decomposition would be algebraically legitimate, but it would not model the emergence of a larger embodiment from one common law core.

The fixed-core condition creates theoretical unity. Each view has the same carrier size, operation type, and complete equational theory.

This permits meaningful comparison among views. It also permits the lift to be interpreted as a coordinated plurality of manifestations of one law.

Exactness matters for the same reason. A quotient that satisfies additional identities is not merely a smaller view of the same theory. It has compressed the source in a way that changes its equational content.

If different coordinates changed the theory differently, the reconstruction would combine theoretically distorted perspectives. The resulting product might recover the carrier while obscuring the intended continuity of law.

By requiring

$Eq(A) = Eq(B)$,

the construction ensures that each quotient core is theoretically complete relative to the lift.

The core is smaller as an object-world, but not poorer as an equational theory.

This is the condition that makes the phrase theory-preserving precise.

8.4 Hypergraph Vertices as Admissible Quotient Policies

For a fixed pair (A,B), consider the set

$Q_B(A)$

of admissible core quotient policies.

At the policy level, an element of $Q_B(A)$ is a congruence θ on A such that

$A/\theta \cong B$.

One may alternatively work with surjective homomorphisms $A \to B$, but several homomorphisms can share the same kernel and differ only by an automorphism of B. For many structural purposes, the kernel is the more intrinsic object.

Each vertex of the reconstruction hypergraph is therefore a lawful identification policy.

The policy answers a binary observational question: which elements of A are treated as equivalent from the perspective of this quotient?

Because the quotient has two elements, the policy divides the carrier into two congruence classes. Yet not every two-block partition is allowed. The partition must be compatible with the operation and induce the selected core type.

The hypergraph vertex set is therefore a filtered subset of all possible binary observations.

This filtering is what gives the hypergraph algebraic meaning. Its vertices are not arbitrary sensors, predicates, or label functions. They are exact core-valued quotient structures.

The number of vertices measures the abundance of lawful perspectives. A lift with few admissible quotients has a rigid observation geometry. A lift with many has greater perspectival redundancy.

But vertex count alone does not determine reconstructibility. What matters is how the corresponding kernels intersect.

8.5 Hyperedges as Minimum Faithful Families

A subset

$E \subseteq Q_B(A)$

forms a reconstruction family when

$\bigcap_{\theta \in E} \theta = \Delta A$.

It is minimal when no proper subset of E has diagonal intersection.

The reconstruction hypergraph has:

vertices equal to admissible core quotient policies;

hyperedges equal to minimal faithful families.

The smallest hyperedge size is the core-view number.

This construction converts the reconstruction problem into a finite combinatorial geometry. Each hyperedge identifies one irreducible way to recover all distinctions in A.

For a two-view lift, the minimum hyperedges are pairs.

For a three-view lift, the minimum hyperedges are triples.

The hypergraph can contain larger minimal families if the structure permits them, although the released four-state atlas focuses primarily on minimum families.

The edge structure reveals which quotient policies are compatible as a complete reconstruction system. Two individually valid views may never appear together in a minimum edge because their kernels overlap too strongly. Other views may be interchangeable.

Thus admissibility is not merely a property of single perspectives. There is also a higher-order admissibility of families.

A lawful view can be valid alone yet insufficient in every pair. Reconstruction is a property of coordinated sets.

8.6 Redundancy, Overlap, and Alternative Reconstructions

The reconstruction hypergraph distinguishes several kinds of structural behavior.

A vertex appearing in every minimum hyperedge is indispensable. Every minimum reconstruction must use that quotient policy.

A vertex appearing in no minimum hyperedge is lawful but inefficient. It may contribute only to larger redundant families.

Two vertices that occur in exactly the same sets of hyperedges are reconstruction substitutes at the combinatorial level.

Disjoint minimum hyperedges represent alternative reconstructions sharing no quotient policy.

Overlapping edges represent reconstructions that share some perspectives while differing in others.

A hypergraph with one minimum edge is rigid. A hypergraph with many edges is redundant.

These distinctions may matter even when the core-view number is identical. Two lifts can both require three views, yet one may admit only one symmetry class of a few triples while another admits many heavily overlapping triples.

The released four-element cases are especially regular because all minimum families within a type lie in one automorphism orbit. Even so, their raw edge counts and overlap patterns may differ.

The hypergraph therefore provides a second layer of classification beneath the view number.

One may think of vB(A) as the minimum cost and the hypergraph as the space of optimal solutions.

8.7 Automorphism Symmetry of the Reconstruction Space

The automorphism group Aut(A) acts on congruences by transport:

$\alpha \cdot \theta$

$\{(\alpha(a),\alpha(b)) : (a,b) \in \theta\}$.

If θ is a core quotient policy, then $\alpha \cdot \theta$ is also a core quotient policy. Intersections are preserved under this action, so minimum faithful families map to minimum faithful families.

The reconstruction hypergraph is therefore an Aut(A)-set with structured incidence.

This action separates labeled multiplicity from intrinsic multiplicity.

Suppose the hypergraph contains six minimum triples. They may represent six genuinely different reconstruction modes, or they may all be images of one triple under automorphisms.

The atlas finds the latter situation in every classified type: all minimum families form one orbit.

This means that the multiplicity of labeled families does not produce multiple intrinsic reconstruction phases at the four-element scale.

The algebra determines one kind of minimum reconstruction, repeated under symmetry.

A stronger invariant than the bare hypergraph would include the full group action:

(H_B(A), Aut(A)).

This pair records both incidence and symmetry.

At larger scales, the orbit decomposition may become especially important. Multiple orbits would indicate qualitatively different minimum reconstruction architectures not related by internal symmetry.

Such a result would suggest that one algebraic world supports several inequivalent modes of lawful observation.

8.8 Beyond a Single Numerical Invariant

A single number is attractive because it permits ordering and comparison. But it necessarily discards structure.

The core-view number does not reveal:

how many admissible quotient policies exist;

how many minimum families exist;

which policies are shared among them;

whether every minimum family is symmetry-equivalent;

whether some policies are indispensable;

how close nonfaithful subfamilies come to reconstruction;

or how the kernels are arranged in the congruence lattice.

The reconstruction hypergraph recovers much of this information, but even it may not be complete.

A richer invariant could weight vertices by quotient-map multiplicity, attach each hyperedge to the resulting diagonal image, record the sizes of all partial kernel intersections, and include the automorphism action.

One could define a reconstruction complex whose simplices are all jointly faithful or partially faithful families.

One could define a separation profile counting how many element pairs remain unresolved by each family size.

One could define a reconstruction polynomial summarizing faithful families by cardinality.

One could define a robustness index measuring how many views can fail before injectivity is lost.

The v0.4.4 artifact establishes the minimum number and the minimum-family hypergraph as the first useful layer. It does not claim that these exhaust the invariant content of lawful reconstruction.

8.9 Partial Reconstruction and Residual Ambiguity

The hypergraph framework naturally supports analysis of incomplete view families.

Given a family F of quotient policies, define its residual ambiguity relation by

$\rho(F) = \bigcap_{\theta \in F} \theta$.

If $\rho(F) = \Delta A$, reconstruction is exact.

If $\rho(F)$ is nontrivial, the quotient $A/\rho(F)$ represents the resolution achieved collectively by F.

This produces a graded view of reconstruction.

With no quotient views, every distinction may remain unresolved.

With one view, A is compressed to the core.

With two views, some lifts are fully reconstructed, while the six cube cases retain residual ambiguity.

With three views, all classified cases can be resolved.

The sequence of intersections describes how objecthood sharpens as perspectives accumulate.

This is important conceptually because reconstruction need not be treated as an all-or-nothing event. Partial families form intermediate object systems. They recover some distinctions while collapsing others.

The residual ambiguity relation may itself be interpreted as the object-resolution policy induced by the available observations.

This links Theory-Preserving Ontogenesis to the broader mathematics of quotient-based object formation. Objects appear at different levels of resolution according to which lawful distinctions are jointly available.

8.10 Reconstruction Hypergraphs and Admissibility

The reconstruction hypergraph can also be described as an admissibility structure.

Each vertex is an admissible observation policy.

Each hyperedge is an admissible jointly complete family.

But not every subset of admissible vertices is itself sufficient.

Thus admissibility has at least two levels:

local admissibility of individual views;

global adequacy of families.

This distinction prevents a common mistake. A collection of individually correct perspectives does not automatically yield a complete account. Their blind spots may coincide.

The kernel intersection makes this exact. If every view identifies the same pair, then multiplying the number of perspectives does not help.

Diversity of lawful views matters because reconstruction depends on complementary distinction.

The hypergraph therefore records not merely quantity of observation but architecture of complementarity.

This may have analogues in distributed sensing, scientific model ensembles, multimodal representation, and collective epistemology. The mathematical content here remains finite and algebraic, but the structural principle is general:

A family succeeds when its lawful limitations intersect only in identity.

8.11 The Core-View Number as Lawful Dimension

The invariant vB(A) may be interpreted as a form of relative dimension.

Ordinary dimension counts the minimum number of coordinates, basis vectors, generators, or independent parameters needed to represent an object.

The core-view number counts the minimum number of B-valued homomorphic coordinates needed to represent A.

This is a dimension constrained by law.

It differs from vector-space dimension because the coordinates are not linear functionals unless the algebras happen to possess linear structure. It differs from set-theoretic coding dimension because arbitrary functions are forbidden. It differs from product rank because every factor is fixed to one exact core.

The phrase lawful dimension captures the main idea:

vB(A) measures how many copies of the core's perspective are needed before the lift becomes fully visible.

For direct-square lifts, the lawful dimension equals the cardinal dimension.

For proper-cube lifts, lawful dimension exceeds cardinal dimension.

The six three-view cases therefore demonstrate dimension inflation under algebraic admissibility.

This interpretation may become increasingly useful at larger sizes, where the gap between unrestricted and lawful dimensions could vary substantially.

8.12 Reconstruction Hypergraphs as Objecthood Signatures

The reconstruction hypergraph may be understood as an objecthood signature of A relative to B.

The lift's elements are not directly placed into the hypergraph. Instead, the hypergraph records the lawful quotient systems through which those elements can be recovered.

This reverses the usual order of description. Rather than beginning with objects and listing their properties, it begins with admissible perspectives and asks which combinations determine the objects.

Two lifts with similar operation tables but different quotient geometries would then have different objecthood signatures.

This perspective aligns with a relational account of identity. An object is characterized by the pattern of lawful distinctions through which it can be separated from all others.

The signature is relative because it depends on the selected core. Objecthood is being measured through one algebraic language of observation.

A future theory might compare reconstruction hypergraphs across different cores, producing a multi-core profile of the same lift. Such a profile would show how objecthood changes under different lawful observational vocabularies.

8.13 Why These Invariants Matter Beyond the Atlas

The exact four-state census is small, but the invariants are not tied to four elements.

For any finite algebras A and B in the same signature, one can ask whether A is residually B through surjective maps and, if so, compute $\nu_B(A)$.

One can construct the associated reconstruction hypergraph.

One can study automorphism orbits.

One can compare the lawful view number with the cardinal lower bound.

One can examine how the invariant behaves under products, subalgebras, quotients, expansions of signature, or changes of core.

This opens a general research program.

Which cores generate lifts with unbounded view overhead?

How does $\nu_B(A \times C)$ relate to $\nu_B(A)$ and $\nu_B(C)$?

Can two equationally equivalent cores yield different view numbers for the same lift?

What varieties force bounded core-view number?

When is the reconstruction hypergraph connected under natural adjacency relations?

Which hypergraphs can arise from finite algebraic reconstruction systems?

These questions move the work beyond one finite catalog while remaining mathematically precise.

The next section turns from the invariants themselves to their conceptual implication: mathematical objecthood can be modeled as lawful reconstruction from coordinated quotient views.

9. Mathematical Objecthood as Lawful Reconstruction

Theory-Preserving Ontogenesis suggests a precise relational account of mathematical objecthood. A mathematical object is not exhausted by a label, a position in a list, or membership in a carrier set. Its identity is secured by the network of lawful distinctions through which it can be separated from every other object while remaining embedded in a coherent algebraic system.

This account does not deny that a finite algebra begins formally with a set of elements. In the implementation, elements receive temporary labels such as 0, 1, 2, and 3. But those labels carry no intrinsic mathematical meaning. They can be permuted without changing the algebra. The mathematical object is therefore not the label itself. It is the invariant role preserved under all structure-respecting relabelings.

The quotient views provide a second route to the same conclusion. No single two-element quotient contains the full four-element object-world. Each collapses distinctions. Yet a jointly faithful family reconstructs the carrier and operation exactly. The objects emerge as unique profiles across the family of lawful views.

The relation can be stated simply:

an element is what no complete lawful family can confuse with another element.

This is stronger than saying that an object is whatever receives a unique code. The code must be generated by homomorphisms that preserve the operation and reproduce the same exact theory. Object identity is therefore relational, operational, and law-governed.

The resulting account is not intended as a universal definition of all mathematical objects. Numbers, spaces, categories, higher types, and infinitary structures involve forms of identity that exceed the present finite setting. The four-state experiment instead provides a minimal exact case in which objecthood is recovered through admissible partial perspectives.

9.1 Objects as Stable Intersections of Views

Let A be a lift and let

$q_1,\ldots,q_m : A \to B$

be a family of lawful quotient maps.

Each quotient map defines an equivalence relation on A. Two elements are equivalent under q_i when that view assigns them the same core value. The object resolution supplied by the full family is the intersection

$\ker(q_1) \cap \ldots \cap \ker(q_m)$.

If this intersection is diagonal, every object is isolated as a singleton equivalence class. No distinct elements remain observationally identical across the entire family.

Objecthood can therefore be represented as the endpoint of successive refinement.

One quotient divides the world into two coarse classes.

A second quotient cross-cuts those classes and may produce finer distinctions.

A third may eliminate the final ambiguity.

The objects are not produced by any one view. They appear as the stable intersections of all required views.

This is analogous to locating a point through coordinates, but with an important difference. Ordinary coordinates can be arbitrary functions. Here each coordinate is constrained by algebraic law. The intersection is therefore not merely geometrical. It is the intersection of lawful equivalence policies.

The phrase stable intersection also emphasizes that identity survives changes of presentation. The quotient maps may be reordered, replaced by automorphically equivalent views, or expressed under different carrier labels. The resulting diagonal relation remains the same.

The individual perspectives vary. The recovered objecthood does not.

9.2 Identity Through Jointly Faithful Observation

A single quotient map cannot faithfully represent a four-element lift in a two-element core. It necessarily identifies distinct elements. Fidelity is therefore not a property of one view but of a coordinated family.

This makes identity distributed.

For an element a in A, its reconstruction profile is

$(q_1(a),\ldots,q_m(a))$.

The family is jointly faithful when distinct elements always have distinct profiles.

Thus identity is encoded through a conjunction of partial observations. No coordinate needs to recognize the whole object independently. The complete identity is carried by the pattern across coordinates.

This differs from an essentialist picture in which every object possesses one intrinsic mark sufficient for identification. In the finite algebraic model, an object can be individuated entirely through its lawful relations to a family of quotient systems.

The word observation should not be interpreted psychologically. The quotient maps are mathematical morphisms, not acts of consciousness. Yet they formalize the structure of a limited perspective: each retains some distinctions and erases others while preserving operation.

Joint faithfulness provides a rigorous criterion for when a set of partial perspectives is sufficient. The criterion does not depend on intuition, visual separability, or descriptive richness. It is exactly the diagonal-kernel condition.

This produces an algebraic principle of identity:

Two elements are the same when every jointly complete lawful observation treats them as the same.

Conversely, distinct elements must be separated by at least one view in every faithful reconstruction family.

9.3 Quotients as Partial but Law-Preserving Perspectives

A quotient is often described as a simplification. It replaces several source elements with one equivalence class and thereby reduces cardinality. But not every simplification is lawful.

The quotient relation must be a congruence. This means that if two elements are treated as equivalent, their behavior under the operation must remain compatible with that identification.

A quotient perspective is therefore selective without being arbitrary. It loses object-level resolution but retains operational coherence.

In Theory-Preserving Ontogenesis, the quotient is even more constrained. It must be isomorphic to the selected core and participate in an exact equality of equational theories.

Each quotient is therefore partial in carrier information but complete in law.

This combination is conceptually important. A view may fail to distinguish objects without being theoretically defective. It can preserve the entire equational theory while compressing the object-world.

The loss is therefore not loss of law. It is loss of resolution.

This separates two forms of incompleteness:

ontic incompleteness, in which not every object distinction is retained;

theoretical incompleteness, in which some laws are lost or altered.

The exact core quotients exhibit the first without the second.

This explains how a smaller structure can remain a complete lawful perspective on a larger one. It does not contain every object, but it contains the full equational theory.

9.4 Ontogenesis Without Addition of New Equational Law

Ordinary growth often suggests the addition of new rules. A larger system may support behaviors unavailable in a smaller system and thereby require a richer theory.

The classified lifts demonstrate a different possibility. Carrier enlargement can occur without equational enlargement.

The lift A contains more elements than the core B, yet

Eq(A) = Eq(B).

No new identity becomes valid merely because the carrier has expanded. No identity of the core fails in the lift.

The lift has a richer finite realization but not a richer equational theory.

This is one of the strongest reasons to call the relation ontogenetic rather than merely representational. Something genuinely new appears at the level of object-world structure:

new elements,

new congruence geometry,

new automorphism behavior,

new quotient configurations,

and new reconstruction hypergraphs.

Yet the theory remains fixed.

The result shows that emergence need not always be identified with new law. New organized worlds can appear as multiple realizations of one law.

This distinction is valuable for the Logos-based program. If law is prior to particular object manifestations, then the same intelligible order may admit several finite embodiments. Ontogenesis becomes variation within law rather than departure from law.

The mathematical experiment does not prove that all natural emergence works this way. It establishes that such emergence is exactly possible in finite algebra.

9.5 Scale Enlargement Without Theoretical Enlargement

The transition from two elements to four is the smallest substantial example of scale enlargement in the artifact.

The core and lift differ in cardinality by a factor of two. The lift may also require two or three quotient perspectives for reconstruction. Its internal structure contains relational information absent from any individual core view.

Nevertheless, the scale increase does not alter the generated variety.

This suggests a distinction between two forms of complexity.

Theoretical complexity concerns which equations hold.

Realizational complexity concerns how a theory is embodied in a particular algebra.

The lift can be more complex in the second sense without being more complex in the first.

A four-element lift may have more subalgebras, more congruences, a larger automorphism group, or a more intricate reconstruction hypergraph than its two-element core. These features are not fully captured by equational theory.

Thus equational equivalence does not erase structural difference. It identifies one level of sameness across different realizations.

This is precisely why theory-preserving ontogenesis is nontrivial. If equality of theory forced isomorphism, no larger lifts could exist. The phenomenon depends on the fact that equational theory is simultaneously strong enough to define law and weak enough to permit multiple nonisomorphic embodiments.

The distinction between theory and realization may generalize beyond finite algebra. A law can remain stable while the number, organization, and resolution of its objects change.

9.6 Object-Worlds as Coherent Relational Systems

The phrase object-world emphasizes that the lift is not merely a collection of isolated objects. Its elements form a closed system under a binary operation.

Every pair of elements produces another element of the same world. The operation creates a complete network of lawful interactions.

An object is therefore defined partly by how it combines with every other object. Two labels that occupy different positions in the operation table are different algebraic objects even if they have no intrinsic descriptions beyond those positions.

The entire world is coherent because the operation is total and closed. Quotient views remain coherent because they preserve that operation. Reconstruction is coherent because the diagonal embedding recovers the operation coordinatewise.

The resulting object-world has three mutually dependent layers:

a carrier of distinguishable elements;

an operation relating every ordered pair;

a theory governing all term-generated relations.

None of these layers is sufficient alone.

A carrier without operation is only a set.

An operation without distinguishable states collapses into triviality.

A list of laws without a realization lacks the finite object-world studied here.

Theory-Preserving Ontogenesis concerns the alignment of all three.

The lift becomes an object-world when its elements, operation, quotient perspectives, and theory form one jointly reconstructible system.

9.7 Relation to Admissibility Algebra and Concept Formation

The construction has a close relationship to the broader idea of admissibility algebra.

An admissibility framework begins by asking which identifications are permitted. An equivalence policy is acceptable only when it respects the relevant structure. In a finite algebra, congruences provide the canonical example: they are precisely the equivalence relations under which the operation descends to a quotient.

The present theory adds a second admissibility condition. A congruence is not selected merely because it defines some quotient. It must define the chosen exact law core.

A third condition appears at the family level. Several individually admissible quotients must be jointly faithful.

Thus object formation proceeds through nested constraints:

a partition must be an equivalence relation;

the equivalence relation must be a congruence;

the quotient must have the required core type;

the core must be equationally exact;

the family of quotients must separate all objects.

This hierarchy makes object formation selective. Most possible binary partitions are rejected. Most arbitrary code systems are rejected. Some lawful quotients remain insufficient in combination.

Concept formation can be described similarly. A concept groups distinct instances under an equivalence policy. The concept is coherent only if the ignored differences do not disrupt the operations or predictions relevant to the task. Several concepts may then interact to individuate an object more finely than any concept alone.

The finite algebraic case therefore provides a rigorous prototype of concept resolution through compatible quotients.

It does not establish a complete theory of human cognition. But it shows how objects can be generated as intersections of admissible abstractions rather than assumed as primitives.

9.8 From Information-Theoretic Bits to Algebraically Admissible Views

The use of two-element cores invites comparison with bits. But the cores in Theory-Preserving Ontogenesis are not ordinary information-theoretic bits.

An ordinary bit has two distinguishable values. Its role is defined primarily by capacity.

An algebraic core has two values and a binary operation. Its role is defined by both capacity and law.

A bit-valued function can assign labels arbitrarily.

A core-valued homomorphism must preserve combination.

This distinction transforms the reconstruction problem.

Under unrestricted binary encoding, every four-element world has dimension two.

Under algebraically admissible encoding, some four-element worlds have dimension three.

The lawful view is therefore a structured bit. It carries not only a state distinction but a rule for how states interact.

One may describe the transition as:

bit → operated bit → homomorphic view → quotient policy → jointly faithful reconstruction.

At each stage, additional constraints are imposed. The informational alphabet becomes part of an algebraic language.

This has implications for how mathematical information might be understood more broadly. Information is often measured independently of the transformations that preserve meaning. But in structured systems, not every encoding is semantically or algebraically valid.

A representation may be compact while destroying the operations that make the represented system intelligible.

Theory-Preserving Ontogenesis therefore suggests that meaningful mathematical compression should be evaluated relative to lawful morphisms. The smallest code is not always the smallest structure-preserving code.

The six three-view worlds provide an exact finite demonstration. They require one more operated bit than unrestricted counting predicts.

9.9 Objecthood as an Invariant Rather Than a Presentation

Every finite implementation must assign labels, order rows and columns, select canonical representatives, and choose one way to display quotient families. These choices are unavoidable but not intrinsic.

The objecthood of the lift lies in what survives them.

Carrier labels can be permuted.

Coordinates can be reordered.

Minimum reconstruction families can be moved by automorphisms.

Core symbols may be exchanged when the core has a nontrivial automorphism.

The invariant content consists of:

the isomorphism class of the lift;

the isomorphism class of the core;

the quotient-policy geometry;

the core-view number;

the reconstruction hypergraph;

and the automorphism action.

This distinction is essential to the philosophical interpretation. Mathematical objects are not the ink marks, array positions, or integer names used to represent them. They are the structures preserved under admissible transformations.

The computational artifact reflects this principle through canonicalization. It removes accidental labeling and retains invariant structure.

In this sense, the methodology itself embodies the relational account of objecthood it studies.

9.10 Lawful Reconstruction and the Limits of Reduction

The lift is reconstructed from core-valued views, but this does not mean it is reducible to any one core.

Each quotient is smaller and theoretically exact, yet none contains the full object structure. Even the family of cores must be coordinated through the pattern of which tuples belong to the diagonal image.

The global lift is therefore more than an unordered collection of core copies. It includes the subalgebra relation selecting the lawful combinations among them.

In a direct square, all coordinate combinations occur.

In a proper cube, only four of eight combinations occur.

The missing combinations are part of the global structure. They cannot be inferred from one coordinate alone.

Thus reconstruction is not simple reduction. The core supplies the local law, while the lift supplies the global coordination.

This produces a balanced relation between parts and whole.

The whole introduces no new equational law.

Yet the whole contains relational organization not present in any part independently.

Theory-Preserving Ontogenesis therefore models a form of emergence that is nonmystical but not trivial. The global world is exactly reconstructible from lawful local views, yet its reconstruction requires their coordinated intersection.

9.11 The Ontological Interpretation and Its Boundary

The mathematical result supports a limited ontological interpretation.

It shows that, in one exact finite setting, objects can be defined through jointly faithful systems of lawful relations. Identity need not be introduced as an unexplained primitive at the level of observation. It can be recovered from the intersection of partial quotient views.

But the result does not establish that all real objects are algebraic quotients, that consciousness constructs reality through congruences, or that physical ontology is exhausted by finite operations.

The mathematics supplies a model, not a universal metaphysical proof.

Its value lies in demonstrating possibility with exactness. A relational account of object formation can be made computationally concrete. Law can constrain observation. Partial perspectives can remain theoretically complete while individually losing identity. Several such perspectives can jointly recover a larger world.

These findings justify further investigation without licensing unrestricted extrapolation.

9.12 From Object Formation to Mathematical Ontogenesis

The full sequence developed in the artifact can now be stated more precisely:

a core carries law;

a quotient view preserves that law while collapsing objects;

a family of views supplies complementary distinctions;

their kernel intersection reaches equality;

the diagonal map reconstructs the larger carrier and operation;

the larger world remains equationally identical to the core.

This is theory-preserving mathematical ontogenesis.

The objects are new as elements of a larger realization.

Their law is not new.

Their identity arises through the coordination of partial lawful perspectives.

Their object-world is recovered as a subdirect pattern inside a power of the core.

The construction therefore gives exact mathematical content to the claim that objecthood may emerge from relation under law.

The next section examines how this result was reached through a distributed human–AI discovery process and what that process implies for mathematical agency, reliability, and participation in intelligibility.

10. Artificial Intelligence and the Distributed Discovery Process

The development of Theory-Preserving Ontogenesis was not the work of one uninterrupted human argument or one continuously persistent artificial intelligence. It emerged across a sequence of conversations, revisions, code packages, audits, objections, failed formulations, corrected definitions, and increasingly precise computational artifacts. The continuity of the work resided not in one memory stream but in the research problem itself.

This distributed origin is important because it mirrors the mathematical theme of the project. No single conversation contained the whole theory. No single model instance preserved every detail of the prior work. Yet the project was reconstructed repeatedly from partial views: manuscripts, definitions, code, release files, numerical outputs, critiques, and the user's persistent governing question.

The discovery process therefore had its own reconstruction geometry. Individual interactions were incomplete, but the accumulated archive was jointly sufficient. Errors introduced in one instance could be exposed by another. Informal claims could be narrowed by code. Code could be checked by frozen outputs. Frozen outputs could be compared with a separate implementation path. The resulting object was not identical to any one conversational contribution.

Artificial intelligence played several roles: translator, formalizer, programmer, critic, editor, and computational intermediary. It did not serve as an autonomous oracle. The human researcher supplied the philosophical direction, rejected insufficient claims, demanded demonstrated novelty, and repeatedly redirected the work away from desired examples toward an operation-first discovery process.

The mathematics supplied a third source of constraint. Once definitions and search spaces were fixed, neither the human investigator nor the AI system could choose the resulting classification freely. The exact numbers, isomorphism types, and minimum reconstruction dimensions emerged from exhaustive finite structure.

The discovery was therefore triadic:

human intention,

machine-mediated formal exploration,

and mathematical resistance.

This triad is central to understanding both the epistemic value and the limitations of the artifact.

10.1 The Human Originating Question

The originating question did not begin as a request to catalog four-element groupoids. It arose from a broader philosophical inquiry into whether mathematics can emerge before its familiar objects and laws have been named.

The human contribution was decisive at this stage. The research direction depended on several choices that were not supplied by the finite algebra itself.

The first was the refusal to begin with a target theorem. The goal was not to program a system to rediscover associativity, Boolean identities, groups, semilattices, or known circuits. The aim was to create a process in which structures and laws could become visible without being selected in advance.

The second was the insistence that emergence from literal nothingness was incoherent. Hardware, software, brains, and formal languages already embody operations and distinctions. This forced a conceptual correction from mathematical creation out of nothing to mathematical disclosure from a lower-level substrate.

The third was the recognition that preserved law should govern the transition from a smaller structure to a larger one. Mere generation of additional elements would not be enough. A larger world had to remain connected to a compact lawful core.

The fourth was the demand for novelty that could be demonstrated rather than asserted. Exploratory programs, suggestive examples, and philosophical claims were repeatedly judged insufficient until they yielded a complete finite classification with exact invariants and reproducible outputs.

These decisions shaped the eventual research question more than any implementation detail did.

The human role was therefore not simply to prompt an AI to produce mathematics. It was to maintain a criterion of significance across many iterations.

10.2 AI-Assisted Formalization Across Multiple Instances

Different AI instances contributed under conditions of incomplete continuity. Each new interaction could inherit documents or summaries, but not necessarily the full internal history of every earlier derivation.

This created both risk and opportunity.

The risk was fragmentation. Definitions might drift. Claims might be strengthened beyond what the computation established. File layouts could be changed without preserving functionality. A later instance might reinterpret a term differently from an earlier one.

The opportunity was repeated reconstruction. Each new instance had to recover the project from explicit artifacts rather than relying solely on conversational familiarity. Weakly documented assumptions became visible because they could not be reconstructed reliably. Important definitions had to be written into source code, mathematical notes, data dictionaries, scope statements, and release manifests.

The distributed process therefore encouraged externalization.

A concept that existed only in one conversational context was fragile.

A concept encoded in definitions, tests, frozen data, and documentation became durable.

This distinction resembles the difference between an informal insight and a formal invariant. The former may depend on the memory of its originator. The latter can be recovered by any competent reader from the artifact.

Multiple AI instances also supplied diverse modes of attention. One might focus on mathematical structure, another on code execution, another on packaging, another on novelty, and another on philosophical interpretation. Their outputs were not automatically consistent, but the inconsistencies themselves became tests of whether the project had been specified clearly enough.

The research archive functioned as the common object reconstructed through these partial perspectives.

10.3 Continuity Through Definitions, Code, and Frozen Results

The project’s continuity does not depend on claiming that every AI instance was the same enduring agent. It depends on the persistence of invariant content.

Several kinds of artifacts carried that continuity.

Definitions preserved the intended mathematical relations among cores, lifts, quotient maps, residuality, and exact theory.

Source code preserved executable interpretations of those definitions.

Tests preserved expected behavior under normal, corrupted, and optimized execution.

Frozen tables preserved the classified structures and summary counts.

Manifests and hashes preserved the identity of the released file set.

Documentation preserved scope, limitations, and interpretation.

Version numbers preserved the history of correction.

These artifacts functioned collectively as an external mathematical memory.

A new AI instance did not need private access to every earlier reasoning step. It could inspect the current definitions, run the code, compare outputs, and reconstruct the state of the project.

This is a stronger form of continuity than conversational recollection alone. A remembered explanation may be mistaken or selectively compressed. An executable release imposes additional constraints.

The frozen outputs are especially important. They prevent the research result from remaining a moving target. A claim such as “there are sixteen lift types” becomes attached to an exact set of rows, operation tables, and invariant fields. A future implementation can agree or disagree with that data.

Continuity therefore resides in reproducibility:

the project remains the same insofar as its formal content can be regenerated and verified.

10.4 Discovery Without Programming a Desired Theorem

One of the principal aims of algebraic abiogenesis was to avoid building known mathematics into the target. A program that is asked to find groups and then tests the group axioms has not discovered the category of groups in the relevant strong sense. It has searched for human-specified instances.

Theory-Preserving Ontogenesis remains partly human-designed, as every formal experiment must. The signature, carrier sizes, quotient conditions, and residual criterion were selected in advance. But the exact classification was not.

The program was not told:

find sixteen lift types;

find a 10/6 division;

find six systems requiring three views;

find five meet-semilattices;

or make every minimum family lie in one automorphism orbit.

Those outcomes arose from the finite search.

This distinguishes target conditions from discovered consequences.

The target condition was theory-preserving residual reconstruction.

The discovered consequences included the exact census and its internal regularities.

No computational search is completely presupposition-free. The meaningful question is whether the answer was encoded in the acceptance criterion. Here it was not. The criteria determined what counted as an admissible object, but they did not determine how many such objects existed or what reconstruction dimensions they would exhibit.

The process therefore achieved a modest but real form of open-ended mathematical discovery within a bounded universe.

It did not discover mathematics from nothing.

It discovered the structure of a precisely defined space without being given its atlas.

10.5 Computational Resistance as Evidence of Mathematical Constraint

A central feature of mathematical discovery is resistance. The investigator may choose definitions and methods, but cannot choose which consequences follow.

This resistance appeared repeatedly during the project.

Broad claims had to be narrowed because the code established only residual reconstruction, not every theory-equivalent extension.

Flat GitHub packaging proved incompatible with activated workflow paths despite explanatory documentation.

Candidate interpretations had to be separated from exact results.

The core-view number could not be reduced to binary coding because six cases violated the cardinal expectation.

The number of lift types was not adjustable without changing definitions or introducing errors.

Such resistance is epistemically valuable. It indicates that the project is not merely rhetorical construction.

The machine itself is not the ultimate source of the resistance. Code can contain mistakes. But when separate implementations, structural theorems, frozen data, and corruption tests converge, the resistance becomes increasingly attributable to the mathematical domain rather than one software path.

This does not mean that computation converts every output into truth. It means that disciplined computation can become an instrument through which finite mathematical necessity is encountered.

The user and AI designed the question.

The classified universe answered it.

10.6 Independent Paths, Audits, and Epistemic Reliability

Computational mathematics requires a layered account of confidence.

A successful run is the weakest layer. It shows only that a program completed.

Agreement with expected summary counts is stronger, but those counts may have been hard-coded or generated by the same faulty logic.

Byte-for-byte reproduction of frozen files is stronger still, because it confirms deterministic regeneration of the artifact. Yet reproduction by the same implementation may reproduce the same bug.

A second implementation path provides another layer. It can use different data structures, control flow, or enumeration order while calculating the same mathematical objects.

Structural reasoning supplies a further layer. The embedding and quotient relations prove equality of equational theories without testing infinitely many identities.

Independent subclass recovery adds another check. The meet-semilattices can be generated separately from their defining laws and compared with the general catalog.

External specialist review would provide an additional social and mathematical layer not yet supplied by internal computation.

These layers should not be collapsed under one word such as verified. Each supports a different claim.

The artifact is internally audited.

Its outputs are reproducible.

Its central theory-preservation argument is mathematically standard.

Its classification has been independently recomputed within the present evaluation environment.

It has not yet received broad external replication or formal proof-assistant certification.

Epistemic reliability comes from making these distinctions explicit rather than exaggerating certainty.

10.7 AI as an Instrument of Participation in Intelligibility

Within the Logos-based interpretation, AI can be understood as an instrument through which finite intelligence extends its participation in mathematical order.

This does not require attributing divine inspiration, consciousness, or spiritual status to the machine. An instrument can participate functionally in discovery without possessing the full form of agency held by its user.

A telescope extends vision without becoming an astronomer.

A proof assistant extends formal checking without becoming the source of logic.

A computer search extends finite enumeration without becoming the ground of mathematical truth.

Large language models add a distinctive capacity. They can move between philosophy, notation, source code, documentation, criticism, and exposition. They mediate among conceptual levels that were previously separated among human specialists and software tools.

This mediation can accelerate research, but it also introduces linguistic confidence without guaranteed correctness. AI-generated mathematics must therefore be subjected to stronger externalization and checking than polished prose alone suggests.

The present project illustrates both sides.

AI helped translate a metaphysical sequence into exact algebraic conditions.

It also produced packaging advice that required later correction.

It helped generate and refine the code.

It also required independent execution and critique.

It provided continuity across a large conceptual program.

It did not possess infallible continuity across instances.

The appropriate model is neither autonomous machine genius nor passive text completion. It is instrumented participation under constraint.

10.8 Authorship, Agency, and the Limits of Machine Illumination

Authorship in AI-assisted research cannot be resolved merely by counting words or code lines. The relevant questions include who originated the problem, who selected the standards of success, who accepted or rejected directions, who assumes responsibility for publication, and who can defend the claims.

In this project, the human author supplied the originating philosophical framework and sustained research intention. The AI systems contributed substantial formal, computational, and textual labor. The mathematical result emerged through their interaction with a constrained finite domain.

The phrase collaboration with AI is therefore accurate, but it should not erase asymmetry.

The human researcher possesses biographical continuity, publication responsibility, and the capacity to place the work within a long-term intellectual program.

The AI instances supplied powerful but discontinuous assistance.

The mathematics supplied consequences attributable to neither party's preference.

Machine illumination must also be distinguished from divine illumination. An AI system can expose patterns, generate conjectures, and organize proofs. These operations do not establish that the machine has participated consciously in the divine light as understood in classical theology.

The theological metaphor should therefore be applied with restraint.

AI may function as a medium through which human inquiry reaches structures previously inaccessible.

It does not follow that the machine possesses spiritual insight.

10.9 Distributed Intelligence and Reconstructed Research Identity

The distributed development process offers a second-order analogy to the mathematical content.

Each AI instance acted as a partial view of the research program.

The user's continuing direction supplied the common core of intention.

The documents, code, and frozen results served as quotient-resistant memory.

No single interaction was sufficient to reconstruct the full project, but the family of artifacts and conversations was jointly faithful.

This analogy should not be made literal. Conversations are not algebra homomorphisms, and research continuity is not a finite congruence lattice. Yet the structural resemblance is useful.

Identity can persist across changing realizations when sufficient invariants are preserved.

A project need not remain in one mind or one session to retain coherence.

It can exist as a distributed object reconstructed through stable relations among records, definitions, and results.

This is increasingly relevant in AI-assisted scholarship. Research may pass among multiple models, tools, collaborators, and repositories. Continuity will depend less on one uninterrupted authorial workspace and more on the quality of externalized invariant structure.

Theory-Preserving Ontogenesis was not merely developed by a distributed process. It also demonstrates why such a process can remain coherent.

10.10 The Human–AI–Mathematics Triangle

The discovery process can be represented as a triangle.

At one vertex is human purpose.

At a second is artificial computational mediation.

At the third is mathematical structure.

Human purpose without computation may remain speculative or too slow for exhaustive search.

Computation without human purpose may generate vast amounts of structure without selecting significant questions.

Human and machine activity without mathematical resistance may produce only internally consistent fabrication.

The research result appears when all three interact.

The human asks a question capable of significance.

The AI translates and explores it at scale.

The mathematics constrains the answer.

This triangle clarifies agency.

The user did not manually enumerate the atlas.

The AI did not originate the full philosophical program independently.

The mathematical result was not created by consensus between user and AI.

Each component made a distinct contribution.

This division also explains why critique is necessary. Errors can enter through human framing, machine implementation, or mathematical misunderstanding. Reliability requires checks appropriate to each source.

10.11 Toward Autonomous Mathematical Discovery

The project remains far from fully autonomous theorem discovery. Human choices still determine the signature, carrier sizes, admissibility conditions, invariants, and interpretation.

Yet it advances beyond target-driven verification in several ways.

It searches all anonymous two-element binary operations.

It constructs larger residual worlds without demanding familiar named structures.

It identifies exact law classes.

It discovers minimum reconstruction numbers.

It produces a complete atlas rather than one requested example.

It recognizes familiar subclasses only after general enumeration.

A more autonomous successor would need to select promising signatures, formulate its own invariants, identify unexpected phase boundaries, propose conjectures from the data, and construct proofs explaining them.

It would also need mechanisms for distinguishing mathematically deep results from exhaustive but unilluminating catalogs.

The present artifact supplies one possible template:

begin with anonymous operations;

define admissible transformations;

enumerate exact finite worlds;

measure reconstruction structure;

detect deviations from naive bounds;

and convert those deviations into general conjectures.

The third-view phenomenon is precisely the kind of deviation that can guide a broader theory.

10.12 The Discovery Process as Progressive Disclosure

The history of the project supports the paper's central philosophical interpretation.

The final result was not visible at the beginning.

The initial Logos-based question concerned the inception of mathematics at a very broad level.

That question led to algebraic abiogenesis.

The substrate objection corrected the claim of emergence from nothing.

The search for target-free novelty led toward law-preserving enlargement.

Universal algebra supplied the exact language.

Computation disclosed the finite atlas.

Audit separated reproducible result from hopeful interpretation.

The process therefore followed a sequence of progressive disclosure:

metaphysical intuition → conceptual correction → formal definition → computational universe → invariant structure → certified result.

At each stage, some possibilities were eliminated and others became visible.

This is not proof that mathematical discovery originates in the Divine Logos. It is an example of how a Logos-based philosophy can generate a disciplined research trajectory rather than remaining a decorative theological interpretation.

The next section evaluates the novelty of the resulting mathematics, distinguishes its new contributions from classical universal algebra, and states the limitations that must accompany any defensible claim.

11. Novelty, Limitations, and Standards of Claim

The mathematical value of Theory-Preserving Ontogenesis depends on distinguishing what is genuinely new from what is inherited from classical universal algebra. The project uses familiar structures: finite algebras, congruences, quotients, direct powers, subdirect embeddings,

varieties, automorphism groups, and exhaustive classification. None of these tools is novel by itself.

The novelty lies in how they are combined, constrained, measured, and interpreted.

The project asks for larger algebras that are simultaneously quotient-linked to a fixed smaller core, residually reconstructible from copies of that core, and exactly equationally equivalent to it. It then classifies the smallest nontrivial case, computes the minimum number of lawful core views required for reconstruction, and records the full space of minimum families through a reconstruction hypergraph.

The resulting contribution is therefore not a new foundation for universal algebra. It is a new finite classification problem within universal algebra, together with a family of invariants and a philosophical interpretation motivated by mathematical inception.

A defensible account of novelty must remain narrower than the broad language of ontogenesis might initially suggest. The code does not classify every theory-preserving extension. It does not prove that all mathematical objects arise through quotient reconstruction. It does not establish a metaphysical doctrine of Logos. It does, however, deliver an exact and reproducible result that appears not to have been previously organized in this form.

The strongest claims are those directly supported by the exhaustive atlas. The weakest are those that extrapolate from the atlas to all of mathematics, cognition, physics, or theology. Maintaining that distinction is essential to the credibility of the work.

11.1 Standard Universal-Algebraic Foundations

The mathematical framework rests on well-established principles.

Homomorphisms preserve term operations.

Kernels of homomorphisms are congruences.

Congruences determine quotient algebras.

Direct products carry operations coordinatewise.

A family of homomorphisms yields an injective diagonal map exactly when its kernels intersect at the diagonal.

An algebra embedded into a product of copies of another algebra satisfies every identity of the factor algebra.

A homomorphic image satisfies every identity of its source.

Birkhoff's theorem characterizes varieties through closure under homomorphic images, subalgebras, and products.

Subdirect representations and residual separation are classical topics.

Minimum factor decompositions and congruence-based reconstruction also have precedents in universal algebra and related areas.

The present project should not describe these as newly discovered facts. They are the inherited machinery that makes the classification possible.

The same caution applies to the observation that two unrestricted binary coordinates can encode four elements. That is elementary information theory.

The novelty begins only when these classical ingredients are placed under the fixed-core and exact-theory constraints and then exhaustively classified.

11.2 Novelty of the Exact Sixteen-Type Atlas

The first principal contribution is the exact atlas of four-element binary algebras residually reconstructible from two-element equational cores.

The atlas is not merely a list of all four-element groupoids. Such catalogs belong to an older tradition of finite-algebra enumeration. The present search is organized by a different relation:

which four-element algebras arise as subdirect lifts of two-element cores while sharing their exact equational theories?

The classification therefore cuts across the ordinary universe of small groupoids according to a residual and theory-preserving criterion.

Its exact result consists of sixteen nonisomorphic core-lift types.

The atlas includes more than the existence of each lift. It records core identity, canonical lift form, quotient maps, minimum faithful families, automorphism behavior, splitting, and reconstruction dimension.

I am not aware of an earlier published classification organized around precisely this combination of criteria and invariants. That supports a claim of probable novelty.

The appropriate wording is:

This work appears to provide the first exact classification of four-element, one-binary-operation algebras residually reconstructible from two-element exact law cores.

The phrase appears to provide is important until specialist literature review and external peer response are complete.

11.3 Novelty of the Ten-Six Reconstruction Division

The second principal contribution is the division of the sixteen lift types into:

ten direct-square reconstructions;

six proper-cube reconstructions.

The result is stronger than a raw count because it identifies a qualitative boundary.

The ten square cases satisfy the unrestricted cardinal minimum. Two lawful core views suffice.

The six cube cases exceed that minimum. Three lawful core views are necessary.

The existence of this gap is the most conceptually significant finding in the atlas. It establishes that structural preservation can increase coordinate requirements beyond information-theoretic sufficiency.

The number six is part of the exact classification, not a selected family of examples. The computation proves that every three-view case in the stated residual universe has been found.

The novelty claim should focus on this precise fact rather than the general statement that algebraic constraints can restrict coding. Structured coding and homomorphic representation are not new ideas. What appears new is the exhaustive identification of the smallest residual exact-law setting in which the lawful reconstruction number exceeds the cardinal bound, together with the complete count of such cases.

A defensible statement is:

The classification identifies exactly six four-state binary algebras whose reconstruction from a two-state exact law core requires three surjective homomorphic views, despite a two-view cardinality bound.

11.4 Novelty of the Fixed-Core View Number

The core-view number formalizes the minimum reconstruction problem:

vB(A)

the least number of surjective B-valued homomorphisms whose kernels intersect diagonally.

The broad concept of minimizing the number of factors in a subdirect representation is not new. Therefore the project should not claim to have invented minimum subdirect decomposition.

The distinctive features of the present invariant are:

the factor algebra B is fixed;

every coordinate must be surjective onto B;

B is an exact equational core of A;

the invariant is interpreted as the number of lawful views required to reconstruct A;

and it is compared explicitly with the cardinal coding bound.

These restrictions produce a specialized relative invariant.

Its value lies less in the abstract existence of a minimum and more in the conceptual separation it creates between unrestricted dimension and lawful dimension.

The invariant also supports future mathematical questions about products, bounds, varieties, and asymptotic growth.

The novelty claim should therefore be calibrated:

The project introduces the fixed-core view number as a specialized constrained subdirect dimension for theory-preserving reconstruction.

This wording acknowledges classical decomposition theory while identifying the new formulation and use.

11.5 Novelty of the Reconstruction Hypergraph

The reconstruction hypergraph extends the numerical invariant by recording all minimum jointly faithful quotient families.

Its vertices are admissible exact-core quotient policies.

Its hyperedges are minimal families with diagonal kernel intersection.

The automorphism group of the lift acts on this hypergraph.

Each ingredient has classical precedents. Congruence systems, minimal covers, hypergraphs, and automorphism actions are well established. The likely novelty lies in their assembly into an invariant specifically designed for fixed-core algebraic reconstruction.

The hypergraph captures information that the view number discards:

alternative reconstructions;

indispensable views;

overlap among minimum families;

symmetry classes;

and redundancy.

The discovery that all minimum families form one automorphism orbit in every four-element classified type is also part of the new finite atlas.

A strong but cautious claim is:

The reconstruction hypergraph provides a new organizing representation for the minimum exact-core quotient families of a lift, together with their automorphism symmetry.

The project should avoid claiming that no mathematically equivalent construction has ever appeared under different terminology. Such a claim would require a much broader literature search than the present artifact itself supplies.

11.6 Residual Lifts Versus All Theory-Equivalent Extensions

The most important limitation is the residual condition.

The program classifies four-element algebras that embed into finite powers of a two-element core through surjective coordinate maps.

This condition gives exact reconstruction and makes the core-view number meaningful. But it does not include every four-element algebra that might share an equational theory with a two-element quotient.

An algebra could have a two-element quotient B and satisfy

Eq(A) = Eq(B)

without being residually B.

In such a case, the available B-valued homomorphisms would fail to separate some pair of elements. No finite family of core views would reconstruct A through the present method.

These nonresidual exact extensions remain outside the classification.

This distinction prevents a major overstatement. The work does not classify all theory-preserving enlargements of two-element cores. It classifies the residually reconstructible portion.

The correct scope statement is therefore indispensable:

The atlas is exhaustive for four-element binary algebras that map onto and embed subdirectly into finite powers of a two-element exact law core.

The residual limitation is not a hidden technicality. It defines the present meaning of reconstruction.

11.7 Limits of Four-Element and Single-Operation Classification

The first artifact uses one binary operation and carriers of size at most four.

This is the smallest setting capable of exhibiting the third-view phenomenon cleanly, but it remains a tiny portion of universal algebra.

Larger carriers may display:

higher core-view numbers;

larger reconstruction overhead;

multiple automorphism orbits of minimum families;

nonsplitting core quotients;

more than one exact core type;

and reconstruction geometries not visible at order four.

Signatures with multiple operations impose additional compatibility conditions. Groups require multiplication, inversion, and identity when represented in the usual signature. Rings contain addition, multiplication, negation, and zero. Lattices contain meet and join. The lawful view number may change when the signature is expanded because a map preserving one operation may fail to preserve another.

Higher-arity operations and relational structures may produce still different reconstruction behavior.

The computational methods used at order four also do not scale automatically. Exhaustive permutation canonicalization and coordinate-family enumeration become expensive quickly.

The first atlas therefore establishes a base case, not a general classification theorem.

11.8 Why the Result Does Not Prove a Logos Metaphysics

The experiment was motivated by a Logos-based philosophy of mathematics, but its computational success does not prove that philosophy.

The existence of theory-preserving residual lifts is a theorem-like fact of finite algebra. It can be accepted by a formalist, Platonist, structuralist, naturalist, or theologian.

The third-view phenomenon shows that algebraic law constrains reconstruction. It does not establish that the ultimate source of algebraic law is divine.

A secular account could interpret the result as one more property of formal systems generated by chosen definitions.

A Platonist could interpret it as discovery of an objective abstract structure.

A structuralist could emphasize the primacy of relations and morphisms.

A Logos-based account interprets the same result as a finite disclosure of intelligibility grounded in divine reason.

The mathematics underdetermines the metaphysics.

This should be stated explicitly because theological overclaiming would weaken both the philosophy and the mathematics. The project gains credibility by allowing the finite result to stand independently of the metaphysical interpretation.

11.9 Why the Result Nevertheless Tests a Logos-Inspired Program

Although the computation does not prove the Logos, it does test whether the Logos-based philosophy can generate productive mathematics.

A philosophical program can be evaluated partly by the questions it makes visible.

The earlier Logos paper proposed a sequence from intelligibility through relation, operation, invariant, object, law, and theory. Theory-Preserving Ontogenesis translates part of that sequence into a finite experiment.

The result is nontrivial.

Larger objects arise from smaller law-bearing cores.

Their identities are recovered through intersecting lawful perspectives.

Their equational theories remain fixed.

Their reconstruction dimension can exceed cardinality.

The philosophical idea therefore did more than decorate an already completed computation. It shaped the choice of problem.

This is the strongest legitimate connection between the metaphysics and the artifact:

The Logos framework generated a mathematical research program whose first exact experiment produced an apparently novel finite classification.

That is evidence of fruitfulness, not proof of metaphysical truth.

11.10 Internal Audit Versus External Validation

The artifact describes itself as internally audited. This language is appropriate.

The second-path audit recomputes the central structures through a separate implementation path. Reproduction scripts compare generated files with frozen outputs. Tests check corruption detection and optimized execution. An additional clean-room recomputation has recovered the central counts.

These facts support strong confidence in the finite classification.

They should not be described as full independent validation in the sociological sense. External validation would require other researchers to inspect the definitions, implement their own computations, or derive formal proofs.

Publication, peer review, and independent replication remain desirable.

The distinction is especially important in AI-assisted work. Multiple AI instances may produce superficially independent code while sharing training biases, assumptions, or inherited formulations.

True robustness comes from diversity of methods and investigators.

The current status can be stated accurately:

internally audited;

reproducible;

independently recomputed within a separate evaluation;

not yet broadly externally replicated.

11.11 Computational Exhaustiveness and Mathematical Proof

An exhaustive finite computation can establish a classification when three conditions are met.

The search domain must be complete.

The acceptance criteria must be mathematically correct.

The implementation must enumerate and test the domain without error.

The artifact addresses the first condition through finite completeness lemmas and coordinate bounds.

It addresses the second through standard universal-algebraic reasoning.

It addresses the third through testing, duplicate implementations, frozen data, and reproducibility.

The resulting evidence is strong, but computational proof remains dependent on software and hardware correctness.

A future formalization in a proof assistant could certify the finite enumeration more strongly. The operation tables, quotient maps, and kernel intersections are small enough that proof objects may be feasible.

Formal verification would not change the numerical result. It would change the epistemic form of the certificate.

The present classification should therefore be described as an exact computer-assisted result rather than a purely hand-derived theorem.

11.12 Standards for Future Claims

Future releases should preserve several standards established here.

Every use of exhaustive should state the search universe.

Every novelty claim should distinguish standard machinery from new classification.

Every invariant should be defined independently of presentation.

Every philosophical interpretation should be separated from the finite mathematical theorem.

Every computational claim should be reproducible from public code and frozen data.

Every audit should state whether it is internal or external.

Every release should preserve versioned outputs so later corrections do not silently alter prior results.

These standards matter because the broader language of mathematical ontogenesis is ambitious. Ambitious interpretation is sustainable only when the underlying finite claims are unusually precise.

11.13 Calibrated Novelty Verdict

The novelty assessment can now be summarized.

The universal-algebraic machinery is standard.

The residual exact-core classification problem appears new in this form.

The sixteen-type atlas is probably novel.

The exact ten-square and six-cube division is a genuine new finite result.

The six cases exceeding the cardinal reconstruction bound are mathematically significant.

The fixed-core view number is a plausible new specialized invariant built from classical subdirect ideas.

The reconstruction hypergraph is a plausible new organizing invariant.

The result is sufficient for a focused computational universal-algebra paper.

It is not yet a general theory of mathematical ontogenesis.

It does not classify all theory-equivalent extensions.

It does not prove a Logos metaphysics.

The strongest defensible statement is therefore:

Theory-Preserving Ontogenesis v0.4.4 provides an exact computer-assisted classification of four-element binary algebras residually reconstructible from two-element equational cores, identifies six cases in which lawful reconstruction requires three binary views despite a two-view cardinality bound, and introduces fixed-core reconstruction invariants that organize the complete finite atlas.

That claim is substantial, precise, and supportable.

The final substantive section turns from this base case to the larger research program it opens: a general science of lawful mathematical emergence.

12. Toward a General Science of Lawful Mathematical Emergence

The four-state classification establishes a base case rather than a completed theory. Its deeper importance lies in showing that lawful mathematical emergence can be studied as an exact research program. A philosophical question concerning the disclosure of intelligibility has been converted into finite definitions, computable invariants, exhaustive catalogs, and reproducible

tests. The next task is to determine which parts of the observed structure persist beyond the smallest universe.

A general science of lawful mathematical emergence would investigate how larger mathematical worlds arise from smaller law-bearing cores, which forms of theory can remain invariant across such enlargement, how many admissible perspectives are required for reconstruction, and what obstructions prevent complete recovery. It would compare unrestricted informational representation with representation constrained by homomorphism, congruence, continuity, computability, symmetry, or other forms of admissibility.

The central principle would remain:

A larger object-world is lawfully emergent when it can be reconstructed from partial views that preserve the relevant theory, while its global organization is not contained in any one view.

This principle is broader than the present finite binary setting, but it must be extended carefully. Each new domain requires a precise meaning of law, view, reconstruction, theory, and preservation. Equational theory is appropriate for universal algebra, but other settings may require logical, categorical, topological, probabilistic, or computational notions of invariance.

The aim should not be to force all mathematical development into one mold. It should be to build a family of exact models showing how objects, theories, and levels of structure can arise through constrained relations.

12.1 Larger Carriers and Higher Core-View Numbers

The most immediate extension is to larger finite carriers.

For a fixed two-element core B and a lift A of size n, the unrestricted information-theoretic lower bound is

ceiling($\log_2 n$).

The lawful minimum is vB(A), the least number of surjective B-valued homomorphisms whose kernels intersect diagonally.

At four elements, the overhead is either zero or one. The natural next questions are:

Can the overhead exceed one?

Can vB(A) grow much faster than $\log_2 |A|$?

Can a fixed core produce a family of lifts with unbounded lawful reconstruction overhead?

These questions move the third-view phenomenon from an isolated finite observation toward asymptotic theory.

A larger carrier permits more complicated congruence geometry. The admissible core kernels may overlap in ways requiring four, five, or many more views. The minimum faithful family becomes a covering problem over a much larger set of element pairs.

For an n-element algebra, there are

$n(n - 1)/2$

unordered pairs to separate.

A crude universal bound can be obtained by choosing one separating homomorphism for each pair, but the true minimum may be far smaller. The gap between this pair count and the cardinal lower bound defines a broad space in which structural constraints can operate.

Larger searches should not begin by enumerating all $n^{(n^2)}$ binary operations. That approach becomes infeasible rapidly. Core-first generation, congruence-guided construction, canonical augmentation, and symmetry pruning will be essential.

The computational objective should be to identify families and theorems rather than merely produce larger catalogs. For example, one might seek conditions on B under which every finite residual lift satisfies

$vB(A) = \text{ceiling}(\log|B| \, |A|)$,

or conditions forcing strict inequality.

The six cube cases show that equality is not automatic. Larger carriers may reveal whole hierarchies of lawful dimension inflation.

12.2 Multiple Operations and Richer Algebraic Signatures

A single binary operation provides the simplest nontrivial setting, but most familiar algebraic theories use richer signatures.

A group may be represented using multiplication, inverse, and identity.

A ring uses addition, multiplication, negation, and constants.

A lattice uses meet and join.

A module combines additive operations with scalar actions.

A category contains objects, morphisms, composition, and identities, although its many-sorted structure exceeds ordinary single-sorted finite algebra.

Expanding the signature changes admissibility. A map that preserves one operation may fail to preserve another. A partition that is a congruence for multiplication may not be a congruence for addition.

Thus adding operations can only restrict the family of lawful quotient views relative to a reduct of the structure. This restriction may increase the core-view number.

Suppose A and B share the same carrier relation in two signatures, one containing only $\star$ and another containing $\star$ together with an additional operation $\diamond$. Every homomorphism in the expanded signature is a homomorphism in the reduct, but not conversely. Therefore the lawful reconstruction number in the expanded signature may be greater.

This creates a precise way to study the cost of theoretical enrichment.

Additional operations make the system more expressive, but they also make lawful observation more demanding.

One may define signature-relative view numbers:

$\nu_B^{\Sigma}(A)$,

where Σ specifies the operations that must be preserved.

Comparing these numbers across signatures would measure how reconstruction complexity changes when more structure is recognized.

This could illuminate a general principle:

The richer the law that a perspective must preserve, the fewer perspectives qualify, and the more perspectives may be required collectively.

The relationship will not always be monotonic in a trivial way because the core and lift themselves may change under signature expansion. Nevertheless, signature sensitivity should become a central invariant of the broader program.

12.3 Nonresidual Theory-Preserving Extensions

The present theory requires residual reconstruction. Every distinct pair of lift elements must be separated by some core-valued homomorphism.

Yet equational equivalence does not imply residuality.

A larger algebra A may have a quotient B and satisfy

$\mathrm{Eq}(A) = \mathrm{Eq}(B)$

while some pair of elements remains indistinguishable under every homomorphism from A to B.

Such an algebra is theory-preserving but not reconstructible from B-valued views.

These cases are not defects in the current classification. They define a different ontogenetic regime.

Residual lifts exhibit transparent emergence: the larger world can be recovered completely from coordinated core perspectives.

Nonresidual exact extensions exhibit hidden emergence: some distinctions in the larger world are invisible to every perspective expressed in the core language.

This suggests an important dichotomy.

Residual emergence is reconstructible.

Nonresidual emergence contains core-invisible distinctions.

The latter may be especially significant philosophically. A larger object-world can embody no new equational law yet possess internal distinctions that cannot be recovered from any family of exact core quotients.

One could measure the invisible residue by the intersection of all core-map kernels:

RadB(A)

$\bigcap \{\ker(h) : h : A \to B \text{ is a homomorphism}\}$.

If this relation is diagonal, A is residually B.

If it is nontrivial, the quotient

A/RadB(A)

is the largest portion of A visible through B-valued homomorphisms.

The remaining distinctions form a hidden layer relative to the core.

A general theory should classify exact extensions according to this residual radical. It should distinguish:

fully reconstructible lifts;

partially reconstructible lifts;

and completely core-opaque enlargements.

This would expand Theory-Preserving Ontogenesis beyond subdirect representation while retaining a rigorous notion of observational visibility.

12.4 Higher-Arity Relations and Categorical Reconstruction

Universal algebra begins with operations, but mathematical structure also includes relations, partial operations, morphisms, and higher compositional systems.

A graph is defined by a binary relation rather than a total binary operation.

A metric space contains a distance function with codomain outside the carrier.

A topological space is organized by a family of open sets.

A category is defined through many objects and typed morphisms.

A higher category or homotopy type contains transformations among transformations.

Extending lawful reconstruction to these domains requires changing the notion of view.

In relational structures, views should be homomorphisms preserving relations.

In topological settings, they may need to be continuous quotient maps.

In categories, they may be functors, fibrations, or localization maps.

In homotopy theory, reconstruction may require descent, gluing, or homotopy limits rather than ordinary kernel intersections.

The diagonal-map principle has categorical analogues. An object may be reconstructed from a family of maps into simpler objects when the induced map into a product, limit, or sheaf of local data is faithful.

The quotient-kernel intersection condition then becomes one instance of a more general descent condition.

This suggests a categorical extension:

A larger structure is lawfully emergent from a core family when it is recovered as a suitable limit of admissible partial images and remains equivalent under the selected theory functor.

The phrase theory functor is intentionally broad. In universal algebra it records equational theory. In topology it might record homotopy type or a selected family of invariants. In logic it might record a theory under an interpretation.

The danger is excessive abstraction. A useful extension must preserve computable examples and explicit obstructions. The finite algebraic case should remain the reference model against which broader categorical formulations are tested.

12.5 Connections to Coding, Observation, and Distributed Representation

The third-view phenomenon has natural connections to coding theory, distributed sensing, tomography, and representation learning.

In ordinary coding, the principal constraint is injectivity or recoverability under noise. Coordinates can often be arbitrary functions.

In lawful coding, each coordinate must also preserve a selected structure.

This changes the optimization problem from minimizing raw code length to minimizing structure-preserving code length.

One can formulate a general constrained encoding problem:

Given structures A and B, find the smallest m such that A admits a faithful embedding into B^m through admissible coordinate maps.

Different notions of admissibility produce different dimensions.

Linear coding requires linear coordinate maps.

Group coding requires homomorphisms.

Topological coding may require continuity.

Probabilistic coding may require preservation of distributions or Markov structure.

Theory-preserving ontogenesis occupies the universal-algebraic instance of this general problem.

Distributed sensing provides another analogy. Each sensor has limited resolution. The system is reconstructible only when the sensors' blind spots are complementary.

Kernel intersections make this complementarity exact.

A family of observations can fail not because each sensor is poor but because they all erase the same distinction.

This principle may have practical implications for multimodal artificial intelligence. Different modalities or encoders may be individually informative while sharing one structural blind spot. Adding more of the same kind of view may not improve identifiability.

A reconstruction hypergraph could represent which combinations of models, sensors, or abstractions jointly recover a target structure.

The mathematical challenge is to prevent analogy from replacing definition. Practical applications require a clearly specified operation, admissible map, and reconstruction criterion.

12.6 Implications for Autonomous Mathematical Discovery

A future autonomous mathematical system should do more than search for instances satisfying human-specified properties. It should identify new structures, formulate useful invariants, detect unexpected boundaries, and explain why those boundaries occur.

Theory-Preserving Ontogenesis suggests one architecture for such discovery.

First, begin with anonymous finite operations or relations rather than familiar object names.

Second, compute exact equivalence classes under a selected notion of theory.

Third, construct admissible extensions and quotients.

Fourth, measure reconstruction complexity.

Fifth, compare lawful minima with naive cardinal, combinatorial, or information-theoretic bounds.

Sixth, isolate the structures producing strict gaps.

Seventh, search for a general theorem explaining the gap.

The third-view phenomenon illustrates the value of this pipeline. The important discovery is not simply that sixteen structures exist. It is that six violate the naive expectation that two binary coordinates suffice.

Autonomous discovery should prioritize such deviations.

A system might rank results according to:

unexpected inequality;

failure of a natural bound;

emergence of a new orbit type;

abrupt change in reconstruction number;

or mismatch between familiar invariants and observed classification.

It could then generate conjectures relating those deviations to congruence geometry, clone structure, or properties of the core.

The present artifact stops before this level of autonomy. Its definitions and interpretive priorities remain human-supplied. But it establishes a finite environment in which an AI can move from enumeration toward theorem formation.

12.7 From Finite Classification to a General Theory of Ontogenesis

A general mathematical theory should identify which aspects of the four-state result are accidental and which are structural.

Several candidate principles emerge.

The first is theory multiplicity:

One equational theory can possess multiple nonisomorphic finite realizations.

The second is lawful reconstruction:

Some realizations can be recovered from quotient copies of a smaller exact core.

The third is dimension inflation:

The minimum lawful view number can exceed the unrestricted cardinal bound.

The fourth is reconstruction geometry:

The full family of minimum views carries invariant structure beyond the minimum count.

The fifth is symmetry of perspective:

Minimum reconstruction systems should be studied modulo automorphisms.

The sixth is visible and hidden emergence:

Residuality determines whether the larger world is completely visible through the core.

A general theory would investigate these principles across varieties.

For a fixed finite core B, one might define the ontogenetic spectrum

O(B)

{(|A|, vB(A)) : A is a finite exact residual lift of B}.

The spectrum records which carrier sizes and view dimensions occur.

One could refine it by counting isomorphism types, hypergraph types, automorphism orbits, and splitting behavior.

One could define the fecundity function

F_B(n)

the number of n-element exact residual lifts of B up to isomorphism.

One could define the maximum lawful overhead

Ω_B(n)

max [vB(A) − ceiling(log|B| n)]

over n-element lifts A.

These functions would transform the program from one catalog into a comparative theory of cores.

Some cores may be ontogenetically sterile.

Some may admit only direct powers.

Some may generate many proper subdirect lifts.

Some may produce unbounded overhead.

The operation carried by the core would then govern an entire ecology of larger lawful worlds.

12.8 The Logos, Mathematical Constraint, and Progressive Disclosure

The broader theological interpretation must remain disciplined.

The mathematics does not identify the Logos with an operation table.

It does not imply that every finite algebra has theological significance.

It does not prove that a divine intelligence selected the sixteen lift types.

The connection lies at the level of intelligibility and discovery.

The Logos framework proposes that mathematical possibility precedes its human articulation. Finite minds participate in that intelligibility by constructing representations that disclose relations they did not freely determine.

Theory-Preserving Ontogenesis provides a small exact example of this pattern.

The question was intentionally framed.

The definitions were chosen.

The code was constructed.

Yet the classification resisted preference.

The resulting 10/6 division was not contained as an explicit target.

The lawful third view appeared because the finite algebraic relations required it.

This is progressive disclosure in a precise sense. The investigator does not receive a completed theory through revelation. A structured question is asked, an admissible search is performed, and an invariant order becomes explicit.

A Logos-based interpretation sees this not as proof, but as participation:

finite reason encounters an intelligibility that exceeds its initial formulation.

The theological significance resides partly in the humility of this encounter. Mathematical structure is not forced to confirm a metaphysical expectation. It may surprise, obstruct, or redirect the inquiry.

The proper posture is therefore neither triumphalism nor reductionism.

The mathematics should be allowed to speak exactly within its scope.

The metaphysics should explain why such exact intelligibility is worth expecting.

12.9 Conclusion: From Illumined Question to Certified Structure

The research began with a philosophical question: how can mathematics begin if every attempt to derive it already presupposes distinction, relation, operation, or mind?

The answer developed in stages.

Mathematics does not arise from absolute nothingness.

It emerges explicitly from substrates that already permit intelligibility.

The correct model is disclosure rather than creation from nonbeing.

A lawful core can carry a complete equational theory.

A larger algebra can embody the same theory on more elements.

Quotient views can preserve law while losing object resolution.

A jointly faithful family can recover the larger world.

The number of required views can exceed the minimum predicted by cardinality.

The four-state classification makes these claims exact.

It finds ten two-element core types and sixteen four-element residual lift types.

Ten lifts are direct squares.

Six are proper cubes requiring three lawful views.

All exact quotient maps split in the classified universe.

All minimum reconstruction families lie in one automorphism orbit within each type.

Five meet-semilattices are recovered both generally and independently.

The central conclusion is not that three is a mystical number or that finite algebra proves theology. It is that lawful reconstruction is a mathematically distinct phenomenon from unrestricted encoding.

Four objects may fit into two bits, yet their algebraic world may require three law-preserving perspectives.

This difference supports a relational account of mathematical objecthood:

objects are not merely labeled units;

they are stable distinctions reconstructed through coherent lawful views.

The computational artifact therefore stands at an intersection.

It is a finite universal-algebra classification.

It is an experiment in AI-assisted mathematical discovery.

It is a model of object formation through admissible relations.

It is the first exact computational consequence of a Logos-based inquiry into mathematical inception.

Its significance should be judged neither by the breadth of its philosophical vocabulary nor by the smallness of its carriers alone. It should be judged by the precision of the bridge it constructs.

A metaphysical intuition became a formal question.

A formal question became an exhaustive search.

An exhaustive search produced an unexpected invariant division.

That division survived audit and independent recomputation.

The path from Divine Logos to mathematical structure remains philosophical at its origin. But at the level reached here, it has become computationally exact.

The next stage is not to enlarge the rhetoric. It is to enlarge the mathematics.

13. References

1. Youvan, Douglas C. The Inception of Mathematics: Divine Logos, Mathematical Platonism, and the Emergence of Number, Relation, and Law. Preprint, July 2026. DOI: 10.13140/RG.2.2.26115.05925.
2. Youvan, Douglas C. Theory-Preserving Ontogenesis: Exact Finite Experiments on the Smallest Quotient Worlds and the Smallest Families of Quotient Views That Preserve a Complete Equational Theory. Version 0.4.4, GitHub repository, July 2026. https://github.com/DougYouvan/theory-preserving-ontogenesis
3. Youvan, Douglas C. Mathematical It from Bit: Algebraic Abiogenesis, Anonymous Finite Operations, and the Automated Birth of Abstract Algebra. Preprint, July 2026. DOI: 10.13140/RG.2.2.25987.26403.
4. Youvan, Douglas C. Mathematical It from Bit: Stable Quotients and the Algebra Before Abstract Algebra. Preprint, July 2026. DOI: 10.13140/RG.2.2.10022.69448.
5. Youvan, Douglas C. How Mathematics Can Be Born: A Public Guide to Algebraic Abiogenesis and Mathematical It from Bit. Preprint, July 2026. DOI: 10.13140/RG.2.2.30600.99847.
6. Youvan, Douglas C. Object Formation and Conceptual Resolution: Stable Quotients, Complete Refinement Lattices, and Functorial Observation. Preprint, July 2026. DOI: 10.13140/RG.2.2.25921.72800.
7. Youvan, Douglas C. How Mathematics Makes Objects: Observation, Identity, and a New Algebra of Concept Formation. Preprint, July 2026. DOI: 10.13140/RG.2.2.13338.81603.
8. Youvan, Douglas C. The Mathematics of Concept Formation: Discovery Calculus, Explanation Geometry, Theory Dynamics, and the Formal Birth of New Objects. Preprint, July 2026. DOI: 10.13140/RG.2.2.23562.43209.
9. Youvan, Douglas C. No-Corpus Artificial Intelligence from Algebraic Abiogenesis: Policy-Governed Object Formation, Theory Ecologies, and Certified Memory. Preprint, July 2026. DOI: 10.13140/RG.2.2.12146.06088.
10. Youvan, Douglas C. Mathematical Monism Within the Logos: Pythagoras, Heraclitus, Plato, Philo, John, Augustine, Boethius, Maximus, Aquinas, Nicholas of Cusa, Kepler, and Leibniz. Preprint, October 2025. DOI: 10.13140/RG.2.2.26577.80489.
11. The Holy Bible. New Revised Standard Version Updated Edition. National Council of Churches, 2021. See Genesis 1 and John 1:1–18.

12. Plato. Timaeus. Translated by Donald J. Zeyl. Hackett Publishing, 2000.

13. Plato. Republic. Translated by G. M. A. Grube and revised by C. D. C. Reeve. Hackett Publishing, 1992.

14. Aristotle. Metaphysics. In The Complete Works of Aristotle, edited by Jonathan Barnes. Princeton University Press, 1984.

15. Philo of Alexandria. On the Creation. In The Works of Philo, translated by C. D. Yonge. Hendrickson Publishers, 1993.

16. Augustine. Confessions. Translated by Henry Chadwick. Oxford University Press, 1991.

17. Augustine. On Free Choice of the Will. Translated by Thomas Williams. Hackett Publishing, 1993.

18. Augustine. The Trinity. Translated by Edmund Hill. New City Press, 1991.

19. Aquinas, Thomas. Summa Theologiae. Blackfriars edition. Cambridge University Press, 1964–1981.

20. Bonaventure. The Journey of the Mind to God. Translated by Philotheus Boehner. Hackett Publishing, 1993.

21. Nicholas of Cusa. On Learned Ignorance. Translated by Jasper Hopkins. Arthur J. Banning Press, 1981.

22. Kepler, Johannes. The Harmony of the World. Translated by E. J. Aiton, A. M. Duncan, and J. V. Field. American Philosophical Society, 1997.

23. Galileo Galilei. The Assayer. In Discoveries and Opinions of Galileo, translated by Stillman Drake. Anchor Books, 1957.

24. Leibniz, Gottfried Wilhelm. Philosophical Essays. Edited and translated by Roger Ariew and Daniel Garber. Hackett Publishing, 1989.

25. Frege, Gottlob. The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number. Translated by J. L. Austin. Northwestern University Press, 1980.

26. Russell, Bertrand. The Principles of Mathematics. Cambridge University Press, 1903.

27. Hilbert, David. “On the Infinite.” In Philosophy of Mathematics: Selected Readings, second edition, edited by Paul Benacerraf and Hilary Putnam. Cambridge University Press, 1983.

28. Brouwer, L. E. J. “Intuitionism and Formalism.” Bulletin of the American Mathematical Society, volume 20, 1913, pages 81–96.

29. Gödel, Kurt. “What Is Cantor’s Continuum Problem?” American Mathematical Monthly, volume 54, 1947, pages 515–525.

30. Benacerraf, Paul. “What Numbers Could Not Be.” Philosophical Review, volume 74, 1965, pages 47–73.

31. Benacerraf, Paul, and Hilary Putnam, editors. Philosophy of Mathematics: Selected Readings. Second edition. Cambridge University Press, 1983.

32. Maddy, Penelope. Realism in Mathematics. Oxford University Press, 1990.

33. Resnik, Michael D. Mathematics as a Science of Patterns. Oxford University Press, 1997.

34. Shapiro, Stewart. Thinking About Mathematics: The Philosophy of Mathematics. Oxford University Press, 2000.

35. Balaguer, Mark. Platonism and Anti-Platonism in Mathematics. Oxford University Press, 1998.

36. Hersh, Reuben. What Is Mathematics, Really? Oxford University Press, 1997.

37. Lakatos, Imre. Proofs and Refutations: The Logic of Mathematical Discovery. Cambridge University Press, 1976.

38. Wigner, Eugene P. “The Unreasonable Effectiveness of Mathematics in the Natural Sciences.” Communications on Pure and Applied Mathematics, volume 13, 1960, pages 1–14. DOI: 10.1002/cpa.3160130102.

39. Penrose, Roger. The Road to Reality: A Complete Guide to the Laws of the Universe. Alfred A. Knopf, 2004.

40. Tegmark, Max. “The Mathematical Universe.” Foundations of Physics, volume 38, 2008, pages 101–150. DOI: 10.1007/s10701-007-9186-9.

41. Dehaene, Stanislas. The Number Sense: How the Mind Creates Mathematics. Revised and updated edition. Oxford University Press, 2011.

42. Butterworth, Brian. The Mathematical Brain. Macmillan, 1999.

43. Lakoff, George, and Rafael E. Núñez. Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being. Basic Books, 2000.

44. Varela, Francisco J., Evan Thompson, and Eleanor Rosch. The Embodied Mind: Cognitive Science and Human Experience. MIT Press, 1991.

45. Birkhoff, Garrett. “On the Structure of Abstract Algebras.” Proceedings of the Cambridge Philosophical Society, volume 31, 1935, pages 433–454.

46. Burris, Stanley, and H. P. Sankappanavar. A Course in Universal Algebra. Springer-Verlag, 1981. https://www.math.uwaterloo.ca/~snburris/htdocs/UALG/univ-algebra.pdf

47. McKenzie, Ralph, George McNulty, and Walter Taylor. Algebras, Lattices, Varieties. Volume I. Wadsworth and Brooks/Cole, 1987.

48. Grätzer, George. Universal Algebra. Second edition. Springer-Verlag, 1979.

49. Bergman, Clifford. Universal Algebra: Fundamentals and Selected Topics. Chapman and Hall/CRC, 2012.

50. Denecke, Klaus, and Shelly L. Wismath. Universal Algebra and Applications in Theoretical Computer Science. Chapman and Hall/CRC, 2002.

51. Freese, Ralph, and Ralph McKenzie. Commutator Theory for Congruence Modular Varieties. Cambridge University Press, 1987.

52. Hobby, David, and Ralph McKenzie. The Structure of Finite Algebras. American Mathematical Society, 1988.

53. Clark, David M., and Brian A. Davey. Natural Dualities for the Working Algebraist. Cambridge University Press, 1998.

54. Davey, B. A., and H. A. Priestley. Introduction to Lattices and Order. Second edition. Cambridge University Press, 2002.

55. Birkhoff, Garrett. Lattice Theory. Third edition. American Mathematical Society, 1967.

56. Berman, Joel, and Stanley Burris. A Computer Study of 3-Element Groupoids. Computational catalog and technical report, 1994. https://www.math.uwaterloo.ca/~snburris/htdocs/MYWORKS/PAPERS/Groupoid_Catalog.pdf

57. Burris, Stanley. “Computers and Universal Algebra: Some Directions.” Computational and theoretical survey. https://www.math.uwaterloo.ca/~snburris/htdocs/MYWORKS/PAPERS/cua.pdf

58. Ganter, Bernhard, and Rudolf Wille. Formal Concept Analysis: Mathematical Foundations. Springer-Verlag, 1999.

59. Shannon, Claude E. “A Mathematical Theory of Communication.” Bell System Technical Journal, volume 27, 1948, pages 379–423 and 623–656.

60. Ashby, W. Ross. An Introduction to Cybernetics. Chapman and Hall, 1956.

61. Eilenberg, Samuel, and Saunders Mac Lane. “General Theory of Natural Equivalences.” Transactions of the American Mathematical Society, volume 58, 1945, pages 231–294.

62. Mac Lane, Saunders. Categories for the Working Mathematician. Second edition. Springer-Verlag, 1998.

63. Lawvere, F. William, and Stephen H. Schanuel. Conceptual Mathematics: A First Introduction to Categories. Second edition. Cambridge University Press, 2009.

64. Awodey, Steve. Category Theory. Second edition. Oxford University Press, 2010.

65. The Univalent Foundations Program. Homotopy Type Theory: Univalent Foundations of Mathematics. Institute for Advanced Study, 2013. https://homotopytypetheory.org/book

66. Baez, John C., and Mike Stay. “Physics, Topology, Logic and Computation: A Rosetta Stone.” In New Structures for Physics, edited by Bob Coecke. Springer, 2011, pages 95–172.

67. Lenat, Douglas B. “Automated Theory Formation in Mathematics.” Proceedings of the Fifth International Joint Conference on Artificial Intelligence, 1977, pages 833–842.

68. Colton, Simon. Automated Theory Formation in Pure Mathematics. Springer-Verlag, 2002.

69. Schmidt, Michael, and Hod Lipson. “Distilling Free-Form Natural Laws from Experimental Data.” Science, volume 324, 2009, pages 81–85. DOI: 10.1126/science.1165893.

70. Davies, Alex, et al. “Advancing Mathematics by Guiding Human Intuition with AI.” Nature, volume 600, 2021, pages 70–74. DOI: 10.1038/s41586-021-04086-x.

71. Trinh, Thang H., Yuhuai Wu, Quoc V. Le, He He, and Thang Luong. “Solving Olympiad Geometry Without Human Demonstrations.” Nature, volume 625, 2024, pages 476–482. DOI: 10.1038/s41586-023-06747-5.

72. Romera-Paredes, Bernardino, et al. “Mathematical Discoveries from Program Search with Large Language Models.” Nature, volume 625, 2024, pages 468–475. DOI: 10.1038/s41586-023-06924-6.

73. de Moura, Leonardo, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. “The Lean Theorem Prover.” Proceedings of the 25th International Conference on Automated Deduction, Lecture Notes in Computer Science, volume 9195, 2015, pages 378–388.

74. Heule, Marijn J. H., Oliver Kullmann, and Victor W. Marek. “Solving and Verifying the Boolean Pythagorean Triples Problem via Cube-and-Conquer.” Proceedings of the International Conference on Theory and Applications of Satisfiability Testing, Lecture Notes in Computer Science, volume 9710, 2016, pages 228–245.

75. Hales, Thomas, et al. “A Formal Proof of the Kepler Conjecture.” Forum of Mathematics, Pi, volume 5, 2017, article e2. DOI: 10.1017/fmp.2017.1.

76. Sandve, Geir Kjetil, et al. “Ten Simple Rules for Reproducible Computational Research.” PLOS Computational Biology, volume 9, number 10, 2013, article e1003285. DOI: 10.1371/journal.pcbi.1003285.

77. Wilson, Greg, et al. “Good Enough Practices in Scientific Computing.” PLOS Computational Biology, volume 13, number 6, 2017, article e1005510. DOI: 10.1371/journal.pcbi.1005510.

78. Peng, Roger D. “Reproducible Research in Computational Science.” Science, volume 334, 2011, pages 1226–1227. DOI: 10.1126/science.1213847.

The author has published over 5,700 AI-assisted papers at: https://www.researchgate.net/profile/Douglas-Youvan/research . Google Scholar has indexed these preprints, and if you want a particular reference, the AI mode of a Google search (Gemini) has efficiently catalogued these preprints.